The Universe Is Not a Movie: Hilbert-Space Harvesting, Generative Spacetime, Quantum Information, and the Physical Limits of Planck-Scale Reality

Douglas C. Youvan

doug@youvan.com www.youvan.ai

June 17, 2026

A literal Planck-scale image of the observable Universe would contain on the order of 10^185 spatial voxels in a single three-dimensional slice and roughly 10^246 four-dimensional cells across cosmic history. These numbers appear to demand an impossible quantum computer, yet they also reveal a deeper conceptual error: spacetime should not be treated as an independently stored movie frame by frame. Quantum theory permits immense, even infinite-dimensional Hilbert spaces, but finite observers can use only states that can be prepared, controlled, distinguished, transformed, and measured under limits of energy, time, precision, causality, and gravity. This paper develops the idea of Hilbert harvesting: extracting finite, physically meaningful answers from an implicit high-dimensional quantum representation without reconstructing every microscopic degree of freedom. It also advances a generative view of spacetime in which initial or boundary conditions, dynamical laws, causal structure, and selected observables replace exhaustive Planck-cell storage. The resulting framework separates formal state-space dimension from operational information, distinguishes addressing from storage and readout, and proposes a quantum spacetime compiler that answers bounded questions rather than attempting to print reality in full.

Keywords: Hilbert-space harvesting, generative spacetime, quantum information, Planck scale, quantum computing, infinite-dimensional Hilbert space, operational Hilbert dimension, quantum gravity, spacetime information, holographic principle, entropy bounds, quantum measurement, quantum simulation, computational complexity, quantum spacetime compiler, physical limits of computation, emergent spacetime, causal structure, information compression, observable Universe.

OpenAI GPT-5.5 Thinking. CC4.0. 111 pages and 80 references.

https://orcid.org/0009-0004-4206-4253

Outline

1. From the Planck Raster to the Information Problem

1.1 The proposal to image all of spacetime

Introduce the thought experiment of dividing the observable Universe into Planck-scale spatial voxels and Planck-time intervals. Explain why this is a useful limiting calculation even though the Planck length is not experimentally established as a literal pixel size.

1.2 The estimated size of the four-dimensional raster

Derive the approximate 10^185 spatial-cell count for one present-scale cosmic slice, the approximately 10^61 Planck-time intervals across cosmic history, and the resulting rectangular-grid estimate of approximately 10^246 spacetime cells.

1.3 Why the largest number may describe the wrong object

Distinguish the size of an exhaustively printed spacetime movie from the intrinsic physical information required to generate a lawful history. Introduce the central claim that the raster is an output specification rather than necessarily a fundamental ontology.

1.4 The Universe as process rather than archive

Contrast a static database model of reality with a dynamical model in which states, laws, constraints, and interactions generate successive physical conditions.

2. What a Planck-Scale Quantum Computer Would Actually Require

2.1 Independent storage at every spacetime cell

Calculate the logical-qubit requirement when each of 10^246 cells contains one independent binary value. Extend the estimate to fields, particles, curvature variables, and local quantum states requiring multiple qubits per cell.

2.2 Storing one spatial slice and evolving it

Show how sequential simulation could reduce active memory from approximately 10^246 to approximately 10^185 local registers while retaining an operation count potentially approaching 10^246 local updates.

2.3 Logical qubits, physical qubits, and error correction

Explain why fault-tolerant computation requires substantial physical-qubit overhead, syndrome extraction, ancilla systems, decoding, and continuously maintained logical states.

2.4 Parallelism does not eliminate total physical cost

Examine the hypothetical use of approximately 10^185 parallel processing elements and explain why parallel execution trades elapsed time for machinery, energy, communication, and synchronization.

3. Infinite-Dimensional Hilbert Space and Its Apparent Promise

3.1 Finite and infinite Hilbert spaces

Review qubits, qudits, harmonic oscillators, quantum fields, occupation-number bases, and separable infinite-dimensional Hilbert spaces.

3.2 Infinite basis ladders

Explain how systems such as oscillators possess formal states |0>, |1>, |2>, ... without implying that every level is physically accessible.

3.3 Continuous amplitudes and apparent infinite precision

Examine whether complex amplitudes, positions, momenta, phases, and field quadratures could encode unlimited information in their real-number values.

3.4 Mathematical existence versus physical accessibility

Establish the difference between a state existing in a theory and a state being preparable, controllable, stable, distinguishable, and measurable by a finite apparatus.

4. The Operational Hilbert Dimension

4.1 Defining the accessible state space

Introduce d_op(E, R, T, epsilon) as the number of states operationally distinguishable under finite energy E, device size R, available time T, and tolerated error epsilon.

4.2 Defining the Hilbert harvest

Define the operational information capacity as

H_op = log2 d_op.

Interpret H_op as the number of usable qubit-equivalents rather than the formal dimension of the complete mathematical state space.

4.3 Energy-limited access to oscillator states

Show how a maximum usable excitation number n_max produces an effective finite dimension approximately equal to n_max + 1 and an information capacity approximately equal to log2(n_max + 1).

4.4 Precision, noise, and state overlap

Discuss how imperfect preparation, decoherence, detector resolution, and nonorthogonality reduce the number of reliably distinguishable states.

4.5 The Hilbert Harvest Principle

Formulate the principle that quantum systems yield useful information through controllable distinguishability, not through formal dimensionality alone.

5. Addressing, Storage, Superposition, and Readout

5.1 Addressing 10^246 possible cells

Show that approximately

log2(10^246) approximately equals 818

qubits are sufficient to label or address one among 10^246 alternatives.

5.2 Why an address is not a database

Explain why an 818-qubit register can identify one cell, or represent one excitation distributed among many cells, but cannot independently store one value for every cell.

5.3 The configuration-space requirement

Demonstrate that N independent binary cells have 2^N possible configurations. For N = 10^246, the corresponding configuration space has dimension

2^(10^246).

5.4 Superposition does not yield simultaneous classical access

Clarify that a superposition can embody many amplitudes while a measurement returns a limited outcome. Separate quantum parallelism from the mistaken idea that one measurement reveals every branch.

5.5 State tomography and repeated preparation

Explain why reconstructing a general quantum state requires many measurement settings and repeated preparations, with demands that grow rapidly with system dimension.

6. The No-Free-Readout Principle

6.1 Internal representation versus exported information

Distinguish the amount of information implicitly embodied by a state from the amount converted into a classical record.

6.2 The output-size barrier

Argue that producing N independently specified classical output bits ultimately requires N registered or transmitted bits, regardless of how compactly the internal computation represents them.

6.3 Why quantum compression cannot print an uncompressed Universe

Explain that quantum computation may calculate selected global properties without enumerating microscopic details, but it cannot provide a complete 10^246-cell classical movie through a much smaller output channel.

6.4 Measurement as an irreversible bottleneck

Analyze the transition from quantum possibilities to finite recorded outcomes and the thermodynamic, statistical, and informational costs associated with measurement.

6.5 The No-Free-Readout Principle

State the principle that no high-dimensional internal representation automatically provides exhaustive classical access to all of its components.

7. Generative Spacetime Instead of Stored Spacetime

7.1 States, laws, and histories

Represent a physical history through an initial or boundary state, a dynamical rule, an interval of evolution, and a selected family of observables.

7.2 The formal generative description

Introduce a schematic representation such as

spacetime history = state + dynamics + constraints + queries.

Contrast this with the explicit raster representation consisting of one stored value per spacetime cell.

7.3 Temporal redundancy

Explain why adjacent moments are correlated by equations of motion and therefore should not be regarded as independent frames.

7.4 Gauge redundancy and equivalent descriptions

Discuss how multiple mathematical configurations may represent the same physical state, making a naive cell count an overestimate of independent physical content.

7.5 Causality and local relevance

Show how only information within an appropriate causal past can influence a specified observation, permitting query-dependent restriction of the computation.

8. Holography, Entropy Bounds, and Spacetime Compression

8.1 Volume counting versus boundary counting

Compare the approximately 10^185 Planck-volume count of a spatial slice with area-based information measures arising in gravitational thermodynamics.

8.2 Black-hole entropy as an information limit

Explain why attempts to concentrate excessive energy or information into a finite region can lead to gravitational collapse and an entropy proportional to horizon area.

8.3 Cosmological information scales

Examine estimates near 10^122 bits associated with cosmic horizons and clarify the assumptions, ambiguities, and limits of such estimates.

8.4 The nominal compression gap

Compare the approximately 10^246-cell spacetime raster with an approximately 10^122-bit horizon-scale information bound, producing a nominal disparity of approximately 10^124.

8.5 Description complexity versus reconstruction complexity

Explain why a short description can generate an extremely long history while complete decompression may still require enormous time, energy, and output capacity.

8.6 The Spacetime Compression Hypothesis

Propose that physically meaningful spacetime information may be encoded through boundary data, correlations, constraints, and generative laws rather than independent microscopic values at every four-dimensional location.

9. The Quantum Spacetime Compiler

9.1 From universal rendering to bounded questions

Replace the goal of printing all spacetime with the goal of answering specific operational questions about defined regions, observables, and resolutions.

9.2 Inputs to the compiler

Specify the required inputs: boundary or initial state, dynamical law, causal domain, resolution, observable, uncertainty tolerance, and resource limit.

9.3 Causal-cone reduction

Describe how the compiler would identify and retain only those degrees of freedom capable of affecting the requested result.

9.4 Adaptive resolution and multiscale representation

Develop a framework in which fine resolution is used only where necessary, while irrelevant regions remain coarse-grained or implicit.

9.5 Symmetry, conservation, and constraint exploitation

Explain how conserved quantities, gauge conditions, locality, tensor structure, and repeated patterns could reduce computation.

9.6 Observable-oriented output

Illustrate outputs such as transition probabilities, correlation functions, integrated curvature, scattering amplitudes, causal connectivity, or the likelihood of a specified macroscopic history.

10. Physical Limits on Hilbert-Space Harvesting

10.1 Finite energy

Examine the energy required to access high oscillator levels, short wavelengths, large occupation numbers, and rapidly changing states.

10.2 Finite time and quantum speed limits

Discuss limits on how rapidly orthogonal or distinguishable states can be prepared and transformed.

10.3 Finite size and communication

Analyze signal-propagation delays, locality, interprocessor communication, synchronization, and the impossibility of instantaneous global control.

10.4 Decoherence and fault tolerance

Explain how environmental coupling, imperfect gates, leakage, correlated errors, and finite-temperature effects restrict accessible quantum complexity.

10.5 Gravitational back-reaction

Consider the point at which the energy and density of the computer alter the spacetime it is intended to simulate.

10.6 Self-reference and the simulator-inside-the-Universe problem

Explore whether a computer contained within the observable Universe could ever maintain a complete independent representation of the Universe that includes the computer itself.

11. What Infinite Hilbert Spaces Can Productively Provide

11.1 Continuous-variable quantum computation

Assess the use of oscillators and field quadratures for naturally continuous problems without claiming infinite precision.

11.2 Bosonic quantum error correction

Explain how many oscillator levels can be used as structured redundancy to protect a smaller number of reliable logical qubits.

11.3 Quantum sensing and metrology

Show how high-dimensional and continuous-variable states can increase sensitivity to phase, displacement, frequency, acceleration, fields, and spacetime perturbations.

11.4 Analog quantum simulation

Describe how one quantum system can naturally emulate another system with related bosonic, field-theoretic, or many-body structure.

11.5 Multimode communication and entanglement

Examine frequency, time, polarization, spatial-mode, and orbital-angular-momentum multiplexing as finite harvests from larger state spaces.

11.6 Structured function representation

Explain how wavefunctions and field states can embody complex distributions implicitly while yielding selected moments, correlations, or samples.

11.7 Query advantage rather than exhaustive knowledge

Argue that the most useful quantum advantage lies in efficiently extracting carefully chosen observables, not in recovering every amplitude or microscopic variable.

12. Toward an Operational Theory of Computable Reality

12.1 The Universe Is Not a Movie thesis

Synthesize the argument that an exhaustive Planck-scale film is neither a necessary representation of reality nor an attainable computational objective.

12.2 Formal infinity and finite observation

Reconcile infinite-dimensional Hilbert spaces with finite energy, finite distinguishability, finite measurement, and finite observer capacity.

12.3 The Generative Spacetime Principle

State that a physical history is most economically represented by states, laws, constraints, causal relationships, and query procedures rather than by independently stored spacetime cells.

12.4 Research questions

Identify open problems concerning operational Hilbert dimension, holographic compression, emergent geometry, quantum error correction, causal simulation, adaptive spacetime resolution, and observable-specific algorithms.

12.5 Criteria for a quantum spacetime compiler

Propose standards for physical validity, resource accounting, uncertainty reporting, causal consistency, gravitational self-consistency, and reproducibility.

12.6 Philosophical implications

Consider whether physical reality is best understood as a completed four-dimensional object, an evolving quantum process, a network of correlations, or a generator of finite observations.

12.7 Final synthesis

Conclude that the infinity of Hilbert space is not an infinite warehouse available to finite observers. Its genuine value is as a generative substrate from which finite systems can harvest bounded, structured, physically meaningful answers.

13. References

This section will contain the complete numbered reference list. It will include foundational and modern sources concerning the Planck scale, quantum computation, infinite-dimensional Hilbert spaces, continuous-variable systems, quantum measurement, accessible classical information, quantum speed limits, black-hole thermodynamics, entropy bounds, holography, quantum gravity, algorithmic information, tensor networks, emergent spacetime, and the physical limits of computation. References will be numbered in order and collected exclusively in this final section.

Art

1. From the Planck Raster to the Information Problem

1.1 The Proposal to Image All of Spacetime

The proposal to image all of spacetime at the Planck scale begins as a counting exercise but rapidly becomes a question about the nature of physical information. Imagine dividing the observable Universe into cubes whose sides are one Planck length, approximately 1.616 x 10^-35 meters. Each cube would be treated as a spatial voxel, analogous to a pixel in a digital image but extended into three dimensions. To include temporal evolution, one would then record the state of every voxel at intervals of one Planck time, approximately 5.391 x 10^-44 seconds. The result would be a four-dimensional raster containing a presumed elementary cell for every spatial location at every moment.

This construction is deliberately extreme. There is no experimental proof that the Planck length is a literal smallest unit of space, that the Planck time is a universal clock interval, or that spacetime is composed of discrete rectangular cells. The Planck units instead identify scales at which quantum mechanics, gravity, and relativistic spacetime are all expected to become simultaneously important. At such scales, classical ideas such as position, distance, duration, and local geometry may lose their conventional meaning.

Nevertheless, the Planck raster is useful because it provides a maximal-resolution thought experiment. It asks what would be required if reality were represented in the most literal possible way: every location, every moment, and every local property explicitly recorded. The calculation therefore establishes an upper-bound style description of exhaustive imaging rather than a confirmed model of fundamental physics.

The phrase “image all of spacetime” also requires clarification. A normal image records a spatial configuration at one moment. A movie records a sequence of images. A complete spacetime image would instead specify an entire four-dimensional history. It would include not only where matter and fields are located now, but also their values throughout the past and, depending on the definition, throughout the future. The result would not merely be a picture of the Universe. It would be a full-resolution record of becoming.

1.2 The Estimated Size of the Four-Dimensional Raster

The observable Universe has a present comoving radius of roughly 4.4 x 10^26 meters. Approximating it as a sphere, its present comoving volume is

V = (4pi/3)R^3.

Using R approximately 4.4 x 10^26 meters gives a volume on the order of

V approximately 3.6 x 10^80 cubic meters.

A Planck-scale spatial voxel has volume

v_P = l_P^3,

where

l_P approximately 1.616 x 10^-35 meters.

Thus,

v_P approximately 4.22 x 10^-105 cubic meters.

The number of Planck-volume voxels in one present-scale spatial slice is therefore approximately

N_space = V / v_P,

which gives

N_space approximately 8.5 x 10^184.

Rounded to an order of magnitude, one three-dimensional Planck-resolution image of the observable Universe would contain approximately

10^185 spatial voxels.

To extend this into time, the age of the Universe is approximately 4.35 x 10^17 seconds. Dividing this duration by the Planck time gives

N_time = T / t_P,

with

t_P approximately 5.391 x 10^-44 seconds.

This produces approximately

N_time approximately 8.1 x 10^60,

or roughly

10^61 temporal intervals.

A rectangular four-dimensional raster constructed by multiplying the present comoving spatial volume by the full number of Planck-time intervals would contain

N_4D = N_space x N_time,

or approximately

N_4D approximately 6.9 x 10^245.

Rounded to an order of magnitude, this becomes

N_4D approximately 10^246 four-dimensional cells.

This number is not a precise physical census. It assumes that the present comoving volume is assigned to every moment in cosmic history. The early Universe had a much smaller physical scale factor, and a calculation based on physical volume at each moment would yield a lower total. The result also depends on the chosen cosmological model, the meaning of the observable boundary, and whether one counts comoving or proper volume. These refinements may change the coefficient and possibly several orders of magnitude, but they do not alter the basic conclusion. The number remains extraordinarily large.

The figure 10^246 should therefore be understood as a deliberately simple upper-order estimate for an exhaustive Planck-resolution spacetime raster of the observable cosmic history.

1.3 Why the Largest Number May Describe the Wrong Object

The immediate temptation is to interpret 10^246 as the amount of fundamental information contained in spacetime. That conclusion does not follow automatically. The number counts cells in a chosen representation. It does not establish that each cell corresponds to an independent physical degree of freedom.

A digital movie may contain billions of pixels across millions of frames, yet its informational content can be much smaller than the raw raster count when successive frames are highly correlated. A static camera view, for example, may repeat nearly identical backgrounds while only a few objects move. Video compression exploits these correlations by storing changes, patterns, and predictive rules rather than independently recording every pixel in every frame.

Physical history is even more constrained. Successive states are related by dynamical laws. Conservation of energy, momentum, charge, and other quantities restricts what can happen. Relativistic causality limits which events can influence one another. Gauge symmetries allow multiple mathematical descriptions of the same physical condition. Quantum entanglement creates correlations that cannot be understood as independent classical values attached to isolated cells. Gravitational constraints further tie local geometry to the distribution of energy and momentum.

The raw raster therefore treats dependent, constrained, and possibly redundant quantities as if they were independent entries in a database. It measures the size of an uncompressed output, not necessarily the minimum information required to specify the physical history.

This distinction may be expressed schematically as

raw history size = number of cells x values per cell,

while a generative description has the form

compressed history = initial state + laws + constraints + boundary conditions.

The first expression describes what must be printed if every local value is explicitly listed. The second describes what may be sufficient to generate or predict the history. The difference between these two descriptions is the difference between storage complexity and generative complexity.

The number 10^246 therefore answers one question: How many entries would appear in a literal Planck-scale four-dimensional database? It does not answer the deeper question: How many independent physical degrees of freedom are required to generate every observable consequence of that history?

1.4 The Universe as Process Rather Than Archive

The raster model encourages a view of the Universe as a completed archive containing a fixed value in every spacetime cell. This perspective resembles the classical block-universe picture, in which past, present, and future are represented as parts of a single four-dimensional geometry. Even within such a view, however, it does not follow that every point carries an independently stored record.

A process-based description begins elsewhere. It treats the Universe as a state undergoing lawful transformation. In quantum mechanics, an initial state evolves under a dynamical operator. In classical field theory, field values change according to differential equations. In general relativity, geometry and matter evolve together under constraint equations and local dynamics. In each case, later states are not arbitrary additions to an archive. They are generated from earlier conditions, boundary data, and governing rules.

This suggests that the complete spacetime history may be more appropriately understood as the output of a physical generator than as a prewritten movie. The generator need not explicitly store every intermediate condition. It may instead produce local states when required by the dynamics. A computational analog would be a program that generates a long sequence from a compact rule. The sequence may contain trillions of symbols, yet the rule producing it may be short.

The distinction becomes decisive when considering quantum computation. A quantum machine intended to reproduce the Universe need not necessarily hold one register for every four-dimensional cell. It might instead represent a quantum state on a spatial slice, evolve that state according to local laws, and calculate only selected observables. The computational burden

would then depend on the structure of the state, the complexity of the dynamics, the requested resolution, and the quantity being measured.

The Planck raster remains useful because it exposes the impossibility of exhaustive classical rendering. Yet its greatest value may be negative: it shows that a cell-by-cell movie is probably the wrong target. The more meaningful objective is to identify the smallest lawful representation from which physically accessible observations can be generated.

The central transition is therefore from counting locations to counting independent information, from storing histories to generating them, and from demanding exhaustive images to asking bounded physical questions. Once this transition is made, the problem is no longer simply how to build a computer with 10^246 memory cells. It becomes a problem of determining what spacetime information actually exists, what portion is physically accessible, and what must be computed rather than explicitly stored.

2. What a Planck-Scale Quantum Computer Would Actually Require

2.1 Independent Storage at Every Spacetime Cell

The most literal computational interpretation of a Planck-scale spacetime image is to assign an independent memory element to every four-dimensional cell. With approximately

N_4D approximately 10^246

cells, even the simplest possible binary representation would require approximately

10^246 logical bits

if every cell stored one classical value, or approximately

10^246 logical qubits

if every cell carried one independently controllable quantum degree of freedom.

This estimate assumes only one binary distinction per cell. A scientifically meaningful local description would require much more. A single cell might need to encode matter content, field amplitudes, local curvature, momentum, charge, spin, particle occupation, gauge variables, or probabilities associated with alternative local states. Even a crude numerical representation of one scalar field could require dozens of bits or qubits per cell. Multiple fields and geometric variables could raise the local requirement to hundreds, thousands, or more.

If q logical qubits are assigned to each cell, then the total logical memory becomes

Q_logical = q x 10^246.

For q = 100,

Q_logical approximately 10^248.

For q = 1000,

Q_logical approximately 10^249.

These figures still underestimate the true burden if the local quantum state is entangled with neighboring cells. Independent local registers do not by themselves describe the complexity of the global state. The full Hilbert space of N qubits has dimension

2^N.

Thus, a register containing 10^246 qubits would support a global state space of dimension

2^(10^246).

A physical quantum computer would not store each amplitude in a separate classical location. The amplitudes would be embodied collectively in the quantum state of the register. Even so, the machine would still need approximately one physical logical degree of freedom for every independently represented local binary variable.

The literal cell-by-cell architecture therefore converts the raw raster directly into a memory requirement. It assumes that spacetime is a database and that the computer must carry a separately addressable state for every entry. This is the strongest and least compressed interpretation of the problem.

2.2 Storing One Spatial Slice and Evolving It

A more economical strategy would avoid storing the entire four-dimensional history at once. Instead, the machine would represent one three-dimensional spatial slice and evolve it through time.

The active memory requirement would then be based on approximately

N_space approximately 10^185

Planck-scale spatial cells.

If each spatial cell required q logical qubits, the active memory would be

Q_slice = q x 10^185.

This is smaller than the full spacetime archive by a factor of approximately

10^61,

corresponding to the number of Planck-time intervals across cosmic history.

The reduction is enormous, but the resulting machine would still be beyond any ordinary meaning of computation. Even with one qubit per spatial cell, it would require approximately 10^185 logical qubits. If each cell required 1000 logical qubits, the active register would approach 10^188 logical qubits.

The time-evolution burden would remain. If every Planck-scale cell must be updated at every Planck-time interval, the total number of local update events would be approximately

10^185 x 10^61 = 10^246.

The memory has been reduced, but the total amount of local dynamical work remains comparable to the original spacetime-cell count.

The actual gate count could be much larger. A single update may require many elementary quantum gates to implement local interactions, enforce constraints, propagate information, correct errors, and represent nonlinear or gauge-dependent dynamics. If g logical gates are required per cell per time step, then the total gate count is approximately

G_total = g x 10^246.

For g = 1000,

G_total approximately 10^249.

The advantage of slice-by-slice simulation is therefore primarily architectural. It replaces a complete archive with a generative process. Earlier states may be discarded, compressed, or retained only at coarse intervals. The machine stores the present computational state rather than the full history.

This strategy resembles how physical systems themselves evolve. The Universe does not appear to carry a separate physical copy of every earlier state within every present location. Current conditions arise from preceding conditions through dynamical evolution. A simulator should exploit the same principle whenever possible.

2.3 Logical Qubits, Physical Qubits, and Error Correction

A logical qubit is an idealized unit of reliable quantum information. A physical qubit is the actual hardware element used to implement it. Real physical qubits are vulnerable to decoherence, control errors, leakage, imperfect measurement, environmental noise, and unwanted coupling. Large quantum computations therefore require quantum error correction.

A logical qubit is typically encoded across many physical components. Additional ancilla qubits are used to detect errors, extract syndromes, verify operations, route information, and support

fault-tolerant gates. The physical-to-logical ratio depends on the hardware platform, physical error rate, error-correcting code, gate architecture, and target logical failure probability.

Let r denote the average number of physical qubits required per logical qubit. Then

Q_physical = r x Q_logical.

For a hypothetical overhead of

r = 10^3,

a full spacetime archive containing one logical qubit per cell would require approximately

Q_physical approximately 10^249.

A one-slice simulator would require approximately

Q_physical approximately 10^188

for one thousand physical qubits per logical qubit.

These estimates omit supporting infrastructure. A quantum computer also needs control electronics, cooling, power delivery, signal routing, classical processors, decoders, sensors, shielding, timing systems, and communication links. In present architectures, the supporting apparatus occupies far more space and consumes far more energy than the qubits themselves.

Fault tolerance becomes increasingly severe as the computation grows. If a calculation contains G logical operations and the entire computation must succeed with high probability, the logical error rate per operation must be sufficiently small that accumulated failure remains controlled.

A rough requirement is

p_logical << 1/G.

For

G approximately 10^246,

this suggests a logical error probability per relevant operation far below

10^-246,

unless the computation possesses special fault-tolerant structure, local recovery, or error-insensitive observables.

This figure should not be interpreted as a practical engineering specification. It reveals the absurdity of treating every elementary update as part of one globally exact calculation. No plausible physical device can be expected to maintain perfect coherence and precision across

such a computation. Any meaningful simulator would need to tolerate approximation, coarse-graining, stochasticity, local error, and uncertainty.

The demand for exact Planck-scale reconstruction is therefore not merely a memory problem. It is a reliability problem of unprecedented severity.

2.4 Parallelism Does Not Eliminate Total Physical Cost

One might try to solve the temporal problem by assigning a separate processing element to every spatial cell. If approximately 10^185 processors operated simultaneously, one Planck-time update could, in principle, be performed in one parallel computational step.

The history would still require approximately

10^61

sequential time steps.

At one update per Planck time, the simulation would advance in real time at best. It would take approximately the age of the Universe to simulate the age of the Universe.

To run faster than real time, each processor would need to perform updates faster than one Planck time, which has no established physical meaning and would conflict with the very resolution being simulated. Alternatively, multiple temporal segments could be computed independently, but only if their boundary conditions were already known. For a causally evolving system, later states generally depend on earlier states and cannot simply be generated independently.

Parallelism also creates communication demands. Each spatial cell interacts with neighboring cells or with fields extending across a region. Information must be exchanged, synchronized, routed, and corrected. Relativity limits the speed of signal propagation. A processor cannot instantly access the state of a distant region. A physically faithful simulation would need to respect causal structure within the simulated system and within the computer itself.

The machinery required for 10^185 processors would also occupy space and contain energy. If each processor had even a Planck-scale footprint, the machine would already approach cosmic dimensions. If each processor required more than Planck-scale volume, as any known physical device would, the hardware would exceed the available volume of the observable Universe by an overwhelming margin.

Energy poses a comparable problem. Every gate operation, measurement, reset, control pulse, and correction process carries physical cost. Reversible quantum operations can avoid some forms of dissipation in principle, but state preparation, stabilization, measurement, cooling, and error correction remain physical processes.

Increasing parallelism therefore does not make the problem disappear. It trades time for hardware, energy, and coordination. The total spacetime volume occupied by the computation may remain enormous even when the elapsed runtime is reduced.

A useful resource measure is not merely the number of qubits or the number of gates, but the computational spacetime volume:

V_comp approximately number of active logical elements x number of sequential logical cycles.

For the one-slice model,

V_comp approximately 10^185 x 10^61 = 10^246

before accounting for local gate depth, error correction, and communication.

This is closely related to the original Planck-cell count. The burden removed from memory reappears as computational history.

The central lesson is that no simple rearrangement of resources rescues the literal raster model. Full storage demands approximately 10^246 local memory elements. Sequential generation reduces memory but leaves approximately 10^246 update events. Massive parallelism reduces elapsed time but demands approximately 10^185 simultaneous processors and enormous communication infrastructure.

The problem is not that current quantum computers are insufficiently advanced. The problem is that the target itself scales as an exhaustive representation of the entire observable cosmic history.

The only plausible route forward is to abandon exhaustive local reconstruction and exploit physical structure. Symmetry, causality, conservation laws, correlations, holographic constraints, coarse-graining, and observable-specific computation must replace the assumption that every cell deserves an independently stored and updated register.

A useful quantum spacetime computer would therefore not be a larger version of a conventional memory array. It would be a selective generator of physically meaningful observables. Its capability would be measured not by how many Planck cells it can print, but by how efficiently it can answer bounded questions about an implicitly represented quantum history.

3. Infinite-Dimensional Hilbert Space and Its Apparent Promise

3.1 Finite and Infinite Hilbert Spaces

Hilbert space is the mathematical setting used to describe quantum states. A finite-dimensional Hilbert space contains a finite number of mutually orthogonal basis states. A qubit, for example, is represented by a two-dimensional Hilbert space with basis states

|0> and |1>.

A system of n qubits occupies a Hilbert space of dimension

dim H = 2^n.

This exponential growth is one reason quantum computation appears so powerful. A register of 300 qubits is described by a state vector with approximately 2^300 amplitudes, a number larger than ordinary classical storage can explicitly track. Yet the system still contains only 300 physical qubits, and a measurement does not reveal all amplitudes simultaneously.

A qudit generalizes the qubit to d distinguishable basis states:

|0>, |1>, ..., |d - 1>.

Its Hilbert space has dimension d, and its ideal classical labeling capacity is

log2 d bits.

Many physical systems are not naturally limited to two or any other finite number of states. A harmonic oscillator, optical mode, electromagnetic cavity, vibrational mode, or bosonic field mode has the occupation-number basis

|0>, |1>, |2>, |3>, ...

with no largest formal value of n. Its Hilbert space is therefore infinite-dimensional.

Quantum fields extend this structure. Each mode of a bosonic field behaves like a harmonic oscillator, and each mode may contain zero, one, two, or arbitrarily many excitations. When all modes are considered, the full state space becomes vastly larger than the Hilbert space of any finite qubit register.

This mathematical infinity appears, at first, to offer a way around the enormous memory estimates developed in the previous sections. If one oscillator possesses infinitely many distinguishable basis states, then perhaps a single physical system could represent all 10^246 spacetime cells. Since infinity exceeds every finite number, the formal state space is large enough.

That conclusion is mathematically correct but physically incomplete. The issue is not whether the state space contains enough abstract states. It is whether the required states can be reached, maintained, transformed, and measured with finite resources.

3.2 Infinite Basis Ladders

Consider a quantum harmonic oscillator with number states

|n>, n = 0, 1, 2, ...

The energy of the nth state is

E_n = hbar omega(n + 1/2),

where omega is the oscillator frequency.

The basis ladder is infinite, but the energy grows without bound as n increases. A device with finite available energy can therefore occupy only a finite effective portion of the ladder. If its maximum usable energy is E_max, the largest accessible occupation number is approximately

n_max approximately E_max/(hbar omega) - 1/2.

The number of accessible number states is then roughly

d_access approximately n_max + 1.

The corresponding ideal information capacity is approximately

I_access = log2(n_max + 1).

This scaling is severe. Doubling the number of available bits requires squaring the number of reliably accessible states. Adding one bit requires approximately doubling the accessible energy range when the states are evenly spaced.

To use one oscillator as a label for 10^246 spacetime cells would require at least

d_access >= 10^246.

The necessary address information would be

log2(10^246) approximately 818 bits.

In principle, an oscillator with accessible levels from |0> through approximately |10^246 - 1> could label one cell among 10^246 possibilities. Its highest level would have energy on the order of

E_max approximately 10^246 hbar omega.

The exact value depends on omega, but the scaling makes the problem clear. Lowering the frequency reduces the level spacing, yet it also makes adjacent levels harder to distinguish within finite time and noise. The energy requirement and the resolution requirement cannot both be made negligible.

More importantly, such a ladder would still provide only one label. It could state which cell was selected. It would not store an independently chosen value for every cell.

To encode one binary value in each of N cells, the system must represent

2^N

possible configurations. For

N = 10^246,

the number of configurations is

2^(10^246).

A single oscillator has infinitely many formal number states and therefore contains a countable basis large enough to label every finite binary string. However, representing all N-bit configurations by distinct number states would require access to excitation numbers extending toward

n approximately 2^(10^246).

The formal possibility survives, but the required energy, control range, and measurement precision become physically meaningless.

The infinite basis ladder therefore does not remove resource costs. It hides them in the level number.

3.3 Continuous Amplitudes and Apparent Infinite Precision

A second apparent route to unlimited information comes from continuous quantum amplitudes. A pure qubit state can be written as

|psi> = a|0> + b|1>,

where a and b are complex numbers satisfying

|a|^2 + |b|^2 = 1.

Because real and complex numbers can be specified with infinitely many digits, it may appear that a single qubit contains infinite information. One could imagine encoding an enormous

database in the decimal expansion of a, b, a relative phase, or an angular coordinate on the Bloch sphere.

The obstacle is readout. A quantum measurement does not return the full value of a or b. It produces a discrete outcome sampled according to probabilities determined by those amplitudes. Reconstructing the amplitudes requires many identically prepared systems and measurements in multiple bases.

If a probability p is estimated from repeated measurements, ordinary statistical uncertainty decreases approximately as

Delta p approximately 1/sqrt(M),

where M is the number of independent samples.

To gain k additional binary digits of precision by straightforward sampling may require a number of repetitions scaling roughly as

M approximately 2^(2k).

Quantum-enhanced estimation can improve some scaling relations under favorable conditions, but it does not permit infinite precision from finite energy, finite time, or one copy of a state.

The same issue applies to continuous-variable systems. Position, momentum, phase, and electromagnetic quadratures are represented by continuous quantities. A wavefunction psi(x) contains a complex amplitude for every value of x. Formally, this appears to be an uncountable collection of numerical information.

Operationally, however, a detector has finite resolution. Noise blurs nearby values. The system has finite energy and bandwidth. Repeated measurements yield samples from a distribution rather than a direct printout of the entire wavefunction.

A real number can mathematically encode an unlimited sequence of bits in its digits. For example, one might define

x = 0.b_1 b_2 b_3 ...

where the binary digits b_i represent a database. But extracting the nth digit requires measuring x to a precision finer than approximately

2^-n.

For n approaching 10^246, the required precision would be incomprehensible. Small perturbations, thermal noise, quantum uncertainty, gravitational effects, and imperfections in preparation would destroy the relevant digits.

The existence of infinitely precise numbers in the mathematical formalism does not imply that nature supplies infinitely precise measuring instruments.

3.4 Mathematical Existence Versus Physical Accessibility

The essential distinction is between the full formal Hilbert space and the physically accessible subspace.

A mathematical model may contain states with arbitrarily high energy, arbitrarily short wavelength, arbitrarily complex entanglement, or arbitrarily fine phase structure. A finite observer cannot necessarily prepare or detect those states.

For a state to function as information, several conditions must be satisfied.

First, it must be preparable. There must be a physically realizable procedure that transforms an available initial condition into the target state.

Second, it must be distinguishable from alternative states. If two states produce nearly identical measurement statistics under every feasible measurement, they cannot serve as reliable separate memory symbols.

Third, it must be stable long enough to be useful. A state that decays before it can be manipulated or read out contributes little practical capacity.

Fourth, it must be controllable. Computation requires not only storage but also transitions, interactions, gates, routing, and conditional operations.

Fifth, it must be measurable. A stored distinction that cannot be converted into an observable outcome is not operational information for the user.

Sixth, all of these requirements must be satisfied with bounded error using finite energy, finite time, finite apparatus, and finite precision.

The physically relevant question is therefore not

dim H = ?

but rather

How many states in H can be reliably prepared, distinguished, transformed, and measured under the available constraints?

An infinite-dimensional Hilbert space can contain an unlimited sequence of ideal states while a laboratory accesses only a finite low-energy sector. This situation is common throughout physics. The complete mathematical structure is larger than the experimentally usable portion.

The same principle applies to quantum field theory. A field permits modes of arbitrarily short wavelength in the formal continuum model. Shorter wavelengths correspond to higher momentum and energy. Attempting to probe ever smaller distances requires ever larger energies. At sufficiently extreme scales, gravitational effects can no longer be ignored, and concentrating energy into a small region may create a horizon rather than reveal finer spatial detail.

Formal ultraviolet degrees of freedom therefore cannot automatically be converted into usable memory.

Infinite-dimensional Hilbert space remains extraordinarily valuable. It permits natural descriptions of oscillators, fields, photons, vibrations, wavefunctions, and continuous observables. It provides a flexible substrate for quantum simulation, sensing, communication, and error-correcting encodings. It allows a finite experiment to select a useful subspace whose dimension may greatly exceed two.

Its promise is not unlimited storage without cost. Its promise is that a physical platform can expose a large, adaptable family of states from which a finite and controllable subset may be harvested.

The apparent paradox can now be resolved. A single oscillator may possess infinitely many formal states. It may therefore contain, in the purely mathematical sense, enough labels to represent any finite spacetime raster. But representing is not the same as accessing, and accessing is not the same as reading.

The burden of the 10^246-cell problem does not vanish. It is transferred into energy, excitation number, preparation complexity, state separation, noise tolerance, measurement precision, and readout time.

Infinite Hilbert space offers unlimited mathematical room. Physics determines how much of that room can be occupied and used.

4. The Operational Hilbert Dimension

4.1 Defining the Accessible State Space

The formal Hilbert space of a quantum system may be finite, countably infinite, or constructed from continuously labeled states. Yet only a limited portion of that space can be used by a real device. The distinction between formal dimension and usable dimension can be made precise by introducing an operational state space.

Let

d_op(E, R, T, epsilon)

denote the maximum number of quantum states that can be reliably used as distinct physical alternatives by a device subject to the following constraints:

E = available energy,

R = characteristic spatial size,

T = available preparation, processing, and measurement time,

epsilon = maximum tolerated error probability.

The quantity d_op is not simply the number of mathematically orthogonal vectors in the full Hilbert space. It counts states that satisfy an operational test. Each state must be physically preparable, sufficiently stable, distinguishable from the others, transformable under available controls, and readable with total error no greater than epsilon.

This definition turns a vague question about infinity into a finite resource problem. A harmonic oscillator may have the formal basis

|0>, |1>, |2>, ...

but a particular laboratory apparatus may reliably access only the first hundred levels. Its formal Hilbert dimension is infinite, while its operational dimension may be closer to

d_op approximately 100.

A more advanced device might access one thousand levels, but only if it can prepare them selectively, prevent leakage, distinguish neighboring energies, and measure the result with adequate fidelity.

The operational dimension depends on the entire experimental system rather than on the isolated quantum object alone. It includes the control fields, detectors, clocks, calibration standards, error-correcting procedures, environmental isolation, and classical processing required to use the states.

Two states that are mathematically distinct but experimentally indistinguishable do not count as separate operational states. Similarly, a high-energy eigenstate that exists in the model but cannot be prepared before the apparatus fails does not increase usable information capacity.

The definition is also task-dependent. A state may be distinguishable for one measurement but useless for another. A set of optical phases may be suitable for sensing a small displacement but unsuitable for long-term digital memory. A large collection of oscillator states may provide

excellent redundancy for protecting one logical qubit without supporting an equally large number of independently readable symbols.

Thus, d_op should be understood as a constrained and purpose-specific measure of accessible quantum structure.

4.2 Defining the Hilbert Harvest

Once the operational dimension has been identified, its information content can be expressed in qubit-equivalent units.

Define the Hilbert harvest as

H_op = log2 d_op.

The quantity H_op measures how many ideal binary distinctions could be represented by the usable state space. If

d_op = 2,

then

H_op = 1.

If

d_op = 256,

then

H_op = 8.

If

d_op = 10^6,

then

H_op approximately 19.9.

This logarithmic relationship is fundamental. A system with one million reliably distinguishable states does not contain one million bits. It contains enough state alternatives to encode approximately twenty bits in a single selection.

The distinction becomes especially important when discussing infinite-dimensional systems. Formal infinity does not imply infinite H_op. As long as energy, time, control, and precision are finite, the number of operationally distinguishable states remains finite.

A useful comparison can be written as

formal capacity = log2(dim H),

operational capacity = log2 d_op.

For an infinite-dimensional Hilbert space, the first expression is formally unbounded. The second remains finite for any finite physical implementation.

The Hilbert harvest therefore measures the usable portion of quantum possibility. It does not count every state permitted by an idealized theory. It counts only those distinctions that can participate in an actual information-processing task.

This also clarifies why a quantum system may be valuable even when H_op is modest. The usefulness of a state space is not determined only by the number of symbols it can store. Quantum systems can exploit coherence, interference, entanglement, squeezing, and structured dynamics. A small operational Hilbert space may solve a particular problem more efficiently than a much larger classical memory.

The Hilbert harvest should therefore be separated into at least three related capacities:

H_store = reliably storable distinctions,

H_control = reliably transformable distinctions,

H_read = reliably extractable distinctions.

In general,

H_read <= H_control <= H_store <= log2(dim H).

A system may possess many stable states but lack precise controls for moving among them. It may support complex coherent evolution but permit only limited measurement. The smallest of these capacities may determine the practical value of the device.

4.3 Energy-Limited Access to Oscillator States

The harmonic oscillator provides the clearest example of how operational dimension emerges from finite resources.

Its energy levels are

E_n = hbar omega(n + 1/2).

Suppose a device can safely supply no more than E_max to the oscillator. The largest accessible excitation number is approximately

n_max approximately E_max/(hbar omega) - 1/2.

If every level from |0> through |n_max> can be prepared and distinguished, then

d_op approximately n_max + 1.

The corresponding Hilbert harvest is

H_op approximately log2(n_max + 1).

This relationship shows that the information gain grows slowly relative to the energy range. To increase H_op by one bit, the number of accessible levels must double. For equally spaced levels, the maximum required energy must also approximately double.

If one wishes to gain 100 bits from a single oscillator through number-state encoding, the required number of distinguishable levels is

d_op = 2^100,

which is approximately

1.27 x 10^30.

The highest required occupation number would therefore be on the order of 10^30. Even before considering control and measurement, the required dynamic range would be extraordinary.

For 818 bits, enough to label one among approximately 10^246 spacetime cells,

d_op approximately 2^818 approximately 10^246.

The oscillator would need access to occupation numbers extending to approximately 10^246. Its formal Hilbert space permits such levels, but finite physical resources do not.

One might attempt to reduce the energy cost by lowering omega. This decreases the spacing between neighboring levels:

Delta E = hbar omega.

However, closely spaced levels require longer observation times or more precise measurements to distinguish. Energy and time are linked by spectral resolution. A detector attempting to distinguish two frequencies separated by Delta omega generally requires an observation duration on the order of

T approximately 1/Delta omega

or longer.

Reducing the level spacing therefore trades energy demand for time demand and susceptibility to noise. Thermal fluctuations, frequency drift, environmental coupling, and control uncertainty can cause neighboring levels to overlap operationally.

The accessible dimension is consequently determined not only by E_max but also by the smallest resolvable energy separation, the coherence time, and the stability of the oscillator.

A more realistic expression would be

d_op = d_op(E_max, Delta E_res, T_coh, epsilon),

where Delta E_res is the minimum resolvable level spacing and T_coh is the coherence time.

Even if the formal energy spectrum is infinite, the number of reliably separable states within a finite energy window is finite.

4.4 Precision, Noise, and State Overlap

Orthogonality in Hilbert space is an ideal mathematical criterion for perfect distinguishability. Real experiments operate with imperfect preparation and noisy measurement, so practical state discrimination is probabilistic.

Suppose a device attempts to encode symbols using states

|psi_1>, |psi_2>, ..., |psi_d>.

If the overlaps

|<psi_i|psi_j>|

are nonzero for i not equal to j, no measurement can always distinguish the states without error. Even when the ideal target states are orthogonal, noise may transform them into mixed states whose probability distributions overlap.

The operational dimension therefore depends on the tolerated error epsilon. A state set usable at

epsilon = 10^-2

may not be usable at

epsilon = 10^-12.

As the required reliability increases, states must be separated more strongly, measured more carefully, or protected by error correction. This usually reduces the number of states that can fit within a fixed energy, bandwidth, or phase-space region.

For a continuous variable x measured with usable range X and resolution Delta x, a crude classical estimate of the number of distinguishable bins is

d approximately X/Delta x.

The corresponding information is

H approximately log2(X/Delta x).

If the resolution improves by a factor of two, the capacity increases by only one bit. Infinite information would require

Delta x -> 0,

which demands unbounded precision.

Quantum uncertainty places additional restrictions. For conjugate observables such as position and momentum,

Delta x Delta p >= hbar/2.

One may reduce uncertainty in x by increasing uncertainty in p, but the phase-space area occupied by the state cannot be compressed without limit. Squeezed states redistribute uncertainty rather than eliminate it.

Noise also accumulates during control. A d-level system generally requires more complicated pulse sequences, more precise calibration, and stronger protection against leakage than a two-level system. As d increases, energy transitions may crowd together, control fields may unintentionally excite neighboring levels, and the detector may confuse nearby outcomes.

The operational dimension can therefore decrease even while the formal number of available levels increases.

An important consequence follows:

More mathematical states do not necessarily produce more usable states.

A device may perform better by selecting a smaller, well-separated code space than by attempting to exploit every accessible level. This is why high-dimensional systems are often used to protect a small logical subspace rather than to maximize raw storage density.

4.5 The Hilbert Harvest Principle

The preceding analysis leads to the Hilbert Harvest Principle:

The usable information content of a quantum system is determined not by the formal dimension of its Hilbert space, but by the number of states that can be physically prepared, preserved, distinguished, transformed, and measured under finite resource constraints.

This principle can be expressed schematically as

usable quantum capacity =
formal state possibilities
filtered by
energy, time, size, control, stability, precision, and readout.

The word “harvest” emphasizes that the full mathematical state space is not automatically available. A finite device extracts a finite operational subset from a larger reservoir of possibilities.

The principle also separates three questions that are often confused.

First:

How large is the formal Hilbert space?

Second:

How large is the accessible low-energy or controllable subspace?

Third:

How much useful information can be extracted from that subspace for a specified task?

These questions may have radically different answers.

An optical field mode has an infinite-dimensional Hilbert space. A laboratory may control only a few dozen effective states. A measurement protocol may extract only one estimated parameter. The formal dimension, operational dimension, and useful output are therefore distinct.

Applied to the Planck-scale spacetime problem, the Hilbert Harvest Principle prevents a misleading conclusion. It is true that an infinite-dimensional Hilbert space contains enough formal states to label every cell in a finite raster. It may even contain enough basis states to label every possible finite binary configuration. But this mathematical inclusion says nothing about whether those states are physically reachable or readable.

To represent one among 10^246 cell addresses requires an operational dimension of at least

$d_{op} \geq 10^{246}$,

or

$H_{op} \geq 818$.

To store one independent bit in each of 10^{246} cells requires a configuration space with dimension

$d_{op} \geq 2^{(10^{246})}$,

or

$H_{op} \geq 10^{246}$.

The difference between these requirements is immense. The first encodes one selection among many alternatives. The second encodes an independent pattern over all alternatives.

Even a formally infinite Hilbert space does not provide the second capability at finite cost. The energy, control precision, preparation complexity, and readout demands scale with the number of operational distinctions, not with the mere existence of abstract basis vectors.

The Hilbert Harvest Principle also provides a positive direction for quantum technology. The objective should not be to exploit infinity indiscriminately. It should be to identify state structures that yield unusually valuable finite harvests.

Examples include:

large error-corrected logical spaces within bosonic modes,

high-dimensional entanglement for communication,

squeezed states for precision measurement,

analog state evolution for quantum simulation,

multimode encodings for multiplexing,

and implicit representations from which selected observables can be efficiently estimated.

The most productive use of an infinite-dimensional system may be to encode one robust logical qubit, one sensitive physical parameter, or one structured quantum field rather than an enormous raw database.

Operational Hilbert dimension therefore offers a better foundation for discussing quantum computational capability than formal dimension alone. It connects abstract quantum theory to the finite conditions under which information becomes physically meaningful.

The central measure is not how many states a theory permits in principle. It is how many distinctions a real system can make reliably and how much useful knowledge can be harvested from them.

5. Addressing, Storage, Superposition, and Readout

5.1 Addressing 10^246 Possible Cells

A register does not need one physical unit for every location it can name. Digital systems use addresses precisely because a compact symbol can select one item from a much larger set.

If a system contains N possible locations, the number of binary digits needed to specify one location is

n_address = ceil(log2 N).

For

N = 10^246,

the address length is

log2(10^246) = 246 log2(10).

Since

log2(10) approximately equals 3.32193,

the result is

log2(10^246) approximately equals 817.2.

A register of 818 qubits therefore has enough computational basis states to label at least 10^246 distinct alternatives:

2^818 approximately equals 1.75 x 10^246.

This result can initially appear astonishing. The complete four-dimensional Planck raster was estimated to contain approximately 10^246 cells, yet only 818 qubits are required to assign a unique binary address to every cell.

There is no contradiction. An address register answers a restricted question:

Which one of the 10^246 cells is being selected?

It does not answer:

What value is stored in each of the 10^246 cells?

The distinction is familiar in ordinary computing. A 64-bit memory address can identify one among approximately 1.84 x 10^19 possible addresses. The address register does not itself contain the contents of all those memory locations. It contains one number specifying where a read or write operation should occur.

Likewise, an 818-qubit address register could represent one Planck-scale spacetime coordinate. Depending on the chosen mapping, the address might encode three spatial coordinates and one temporal coordinate. If the four coordinates were allocated roughly evenly, each would receive approximately 204 or 205 address qubits.

The register could occupy a basis state such as

|x, y, z, t>,

where the combined binary strings identify one cell in the four-dimensional raster.

The same register could also occupy a superposition:

|psi> = sum_i a_i |i>,

where i ranges over the possible cell addresses. In that state, the register is not assigned to one definite cell in the computational basis. It carries amplitudes associated with many possible addresses.

This is useful for quantum search, interference, amplitude estimation, quantum walks, and many other algorithms. It does not convert the address register into a complete spacetime memory.

The compactness comes from representing one selection among many possibilities. The number of address bits grows logarithmically with the number of possible locations. The amount of data stored independently at those locations grows linearly with their number.

Confusing those two scaling laws is one of the main sources of exaggerated claims about quantum memory.

5.2 Why an Address Is Not a Database

Suppose every spacetime cell contains one binary value

s_i in {0, 1},

where

i = 1, 2, ..., N

and

N = 10^246.

A complete database must specify the string

s_1 s_2 s_3 ... s_N.

The address i identifies one position in this string. It does not determine the value s_i unless some additional memory, function, rule, or physical process supplies that value.

A complete query system therefore requires at least two conceptual components:

an address register |i>,

and a data mechanism returning s_i.

A quantum oracle is often written schematically as

O_s |i>|b> = |i>|b XOR s_i>.

This transformation appears compact, but the oracle must somehow embody the entire function

i -> s_i.

If the values s_i are arbitrary and independent, the complexity has not disappeared. It has been moved into the physical construction or algorithmic description of O_s.

For a structured database, the oracle may be simple. If

s_i = 1

whenever i is even, then a short circuit can calculate the answer from the final address bit. There is no need to store N independent values.

If the database contains a highly compressible pattern, a compact generating program may reproduce it. The memory cost is then determined by the complexity of the rule, not by the number of addresses.

But if the values are algorithmically arbitrary, no comparably short rule exists. The oracle must contain information equivalent to the database itself.

This establishes an important distinction:

address complexity = log2 N,

arbitrary data complexity = N bits.

For

N = 10^246,

these become

address complexity approximately equals 818 bits,

arbitrary binary data complexity approximately equals 10^246 bits.

The difference is not a modest engineering factor. It is the difference between selecting one possibility and specifying an entire pattern.

A library catalog provides an analogy. A short call number may identify one book among millions. The call number does not contain the text of every book in the library. A coordinate can identify one point on Earth, but it does not contain a complete image of the planet. A web address can locate a large file, but it is not the file.

The same principle applies to quantum addresses. The ability to place an address register into a superposition of all locations does not cause the register to contain the independent data stored at all those locations.

It provides coherent access to a set of alternatives, assuming a data source or generative oracle already exists.

5.3 The Configuration-Space Requirement

The storage requirement becomes clearer when expressed in terms of the number of possible global configurations.

For one binary cell, there are two possibilities:

0 or 1.

For two binary cells, there are four possibilities:

00, 01, 10, 11.

For N binary cells, the number of possible configurations is

C_N = 2^N.

A memory capable of representing any possible N-bit pattern must therefore possess at least 2^N reliably distinguishable global states.

For

N = 10^246,

the required state-space dimension is

dim H_required = 2^(10^246).

A register of N qubits supplies exactly this dimension because

dim H_N = 2^N.

Thus, one qubit per independently represented binary cell is not an arbitrary engineering convention. It reflects the number of distinct global alternatives that must be supported.

By contrast, an 818-qubit register has dimension

2^818 approximately equals 10^246.

That is enough for one of 10^246 mutually exclusive labels. It is not enough for one of

2^(10^246)

possible spacetime patterns.

The number of qubits required to represent an arbitrary state selected from M orthogonal alternatives is

n = log2 M.

If

M = 2^(10^246),

then

n = log2(2^(10^246)) = 10^246.

The configuration-space argument therefore reproduces the direct memory count.

This distinction can be stated as follows:

A system with N possible locations needs log2 N address qubits.

A system with one independent qubit-valued degree of freedom at every location needs N qubits.

The difference arises because the first system chooses one location, while the second independently chooses a value at every location.

A quantum field illustrates the second case. The field is not merely located at one lattice point. It assigns local degrees of freedom across the lattice. If each site has d local basis states, then a lattice of N sites has Hilbert-space dimension

dim H = d^N.

The number of qubit-equivalents is

log2(dim H) = N log2 d.

For binary local systems, d = 2, giving N qubits. For larger local dimensions, the requirement increases proportionally to log2 d.

Real quantum fields can have infinite-dimensional local spaces, but practical lattice simulations impose cutoffs on field amplitude, energy, occupation number, or mode count. Without such truncations, the formal global Hilbert space is not computationally manageable.

The configuration-space requirement also explains why entanglement is powerful but not free. Entanglement allows the global state to contain correlations that cannot be decomposed into independent local states. Yet the entangled system still requires physical subsystems or modes capable of supporting the relevant joint Hilbert space.

Entanglement reorganizes information across the system. It does not eliminate the dimensional requirement for representing arbitrary configurations.

5.4 Superposition Does Not Yield Simultaneous Classical Access

A quantum register can be placed into a superposition over many addresses:

|psi> = sum_i a_i |i>.

For N possible addresses, the state may contain N complex amplitudes. This is often described informally as the quantum computer evaluating or representing many possibilities at once.

The phrase is useful only when handled carefully.

Quantum evolution acts linearly on the entire superposition. An operation may transform every amplitude coherently. Different computational paths can interfere, causing some amplitudes to increase and others to cancel. A well-designed algorithm uses this interference to make a desired global property more likely to appear in measurement.

However, measurement does not return the list

a_1, a_2, ..., a_N.

A measurement in the address basis returns one outcome i, sampled with probability

p_i = |a_i|^2.

After measurement, the state is generally altered. The observer obtains one address, not all addresses and not all amplitudes.

Repeated preparation and measurement can estimate the probability distribution, but reconstructing a general distribution over N outcomes requires a number of samples that grows with N and with the desired precision.

Superposition therefore provides coherent processing capacity, not unrestricted parallel readout.

This distinction can be illustrated by preparing the uniform state

|u> = (1/sqrt(N)) sum_i |i>.

The state contains equal amplitude on all N addresses. A measurement yields one randomly selected address. It does not print the entire list of addresses, even though every address participated in the quantum state.

Suppose a quantum oracle attaches data values to the addresses:

sum_i a_i |i>|0> -> sum_i a_i |i>|s_i>.

The output state may correlate each address with its associated data. Yet measuring it still yields one address-data pair in a single run.

Quantum algorithms obtain advantage by manipulating the correlations and phases before measurement. They are designed so that a small number of measurements reveals a useful property such as:

whether a marked item exists,

the period of a function,

an estimated expectation value,

a global parity,

a spectral property,

or a transition amplitude.

They do not generally reveal every individual function value.

The phrase “quantum parallelism” should therefore be interpreted as coherent global transformation under restricted readout. It is not equivalent to having N classical processors each return an independently readable answer.

For the spacetime problem, an 818-qubit address register could be placed into a superposition over all 10^246 cells. A suitable quantum operation might calculate a field-dependent phase, identify cells satisfying a condition, or estimate a global correlation.

But a single measurement could not provide the local state of all 10^246 cells.

The quantum computer may answer a question about the whole raster without rendering the whole raster. That is the useful possibility. The attempt to extract every local value destroys the compression advantage because the output request itself is enormous.

5.5 State Tomography and Repeated Preparation

The difficulty of full readout becomes especially clear in quantum state tomography.

A pure state of a d-dimensional system can be written as

|psi> = sum_j a_j |j>.

The amplitudes a_j are complex. After accounting for normalization and an unobservable global phase, a general pure state requires approximately

2d - 2

independent real parameters.

A general mixed state is represented by a d x d density matrix. After Hermiticity and unit-trace constraints are included, it contains approximately

d^2 - 1

independent real parameters.

For n qubits,

d = 2^n.

A general pure state therefore requires a number of parameters proportional to

2^n,

while a general mixed state requires a number proportional to

4^n.

This exponential parameter count does not mean the quantum computer stores those numbers as explicit classical entries. It means that full classical reconstruction of an arbitrary quantum state is exponentially demanding.

For an 818-qubit address register,

d = 2^818 approximately equals 10^246.

A general pure state over the spacetime addresses would therefore involve approximately 10^246 complex amplitudes. Full reconstruction would require an enormous number of distinct measurements and repeated state preparations.

For a hypothetical register of 10^246 qubits representing one independent qubit per spacetime cell, the Hilbert-space dimension would be

d = 2^(10^246).

A general state would be so large that even writing the number of amplitudes requires an exponential expression.

Tomography becomes impossible long before such scales. The machine would need to prepare the same state repeatedly, measure it in many bases, collect statistics, control systematic error, and store the resulting classical data.

Quantum systems are therefore not transparent containers. Their internal state is known through limited observables and statistical inference.

This opacity is not merely a technological inconvenience. It is part of quantum mechanics. Noncommuting observables cannot all be assigned simultaneously sharp values. Measurement choices matter. Unknown quantum states cannot be copied perfectly for unlimited examination. A single specimen does not reveal its complete wavefunction.

For practical quantum computing, full tomography is usually unnecessary. Engineers characterize selected gates, observables, error rates, and benchmark tasks. Algorithms are designed to produce a small classical output whose correctness can be estimated without reconstructing the entire internal state.

The same strategy must govern any quantum treatment of spacetime. A useful machine would not attempt to reveal its complete universal wavefunction. It would estimate defined observables associated with bounded questions.

The central lesson of this section is that four quantities must remain separate:

address size,

data size,

Hilbert-space dimension,

and readable output.

For approximately 10^246 spacetime cells:

the address requires about 818 qubits,

one independent binary value per cell requires about 10^246 qubits,

the corresponding configuration space has dimension 2^(10^246),

and exhaustive classical readout requires approximately 10^246 output bits even before richer local physics is included.

Superposition can process many alternatives coherently. Entanglement can encode complex correlations. Infinite-dimensional systems can supply vast formal state spaces. None of these removes the need to distinguish between selecting one possibility, embodying an entire configuration, and exporting that configuration as classical information.

This distinction prepares the ground for a broader principle: there is no free readout. A compact or implicit quantum representation may support efficient answers to carefully chosen questions, but it cannot provide an exhaustive classical description without paying the corresponding output cost.

6. The No-Free-Readout Principle

6.1 Internal Representation Versus Exported Information

A quantum system may embody a state whose mathematical description involves an enormous number of amplitudes. That internal representational richness does not imply that an observer can extract an equally enormous classical description from a small number of measurements.

This distinction is essential. The state of a quantum system and the record available to an external observer are not the same object.

An n-qubit pure state can be written as

|psi> = sum_x a_x |x>,

where x ranges over 2^n computational basis states. The state is therefore associated with as many as 2^n complex amplitudes. Yet a measurement in the computational basis returns only one n-bit outcome per run.

The exponentially large state space permits interference among alternatives. It enables transformations that may reveal certain global properties efficiently. It does not provide direct access to every amplitude.

A useful separation is

internal state complexity != classical output size.

The internal system may be highly structured, entangled, or distributed over a large basis. The classical output may contain only one bit, one integer, one sampled configuration, or one estimated expectation value.

This is not a defect of quantum computation. It is part of its power. Quantum algorithms are useful precisely when they transform an implicit high-dimensional state so that a small measurement record contains the answer to a selected question.

The difficulty appears when the requested answer is not small.

If the goal is to recover one binary value for each of 10^246 spacetime cells, then the desired classical record already contains approximately 10^246 bits. The computer may represent the underlying history compactly or implicitly, but the final requested product is still enormous.

The output demand is independent of whether the internal representation is classical, quantum, analog, holographic, compressed, or generative.

A compact rule may generate a long string, but printing the string still requires producing every symbol.

6.2 The Output-Size Barrier

Suppose a device is required to output an N-bit string

s = s_1 s_2 ... s_N.

If the bits are independently specified and the observer must receive all of them, then the output record must contain N binary distinctions.

The machine must therefore perform enough physical interactions to establish those distinctions in an external medium. It may transmit them through a channel, write them into memory, expose them on a display, or encode them into matter. Regardless of the method, the result must support at least

2^N

possible distinguishable records.

This requirement may be called the output-size barrier.

For the Planck-scale spacetime raster,

N approximately 10^246.

Even if the complete history were compressed internally into a much smaller quantum state, an exhaustive classical rendering would still require a record with approximately 10^246 local entries for one binary property per cell.

If each cell carried b classical bits, then

B_output = b x 10^246.

For a modest 32-bit value per cell,

B_output approximately 3.2 x 10^247 bits.

If the description included many local fields, uncertainties, tensor components, particle data, or quantum-state parameters, the output would be far larger.

The output-size barrier is not merely a limitation on computer memory. It constrains communication, observation, and physical recording.

At a transmission rate of r bits per second, the minimum output time is

T_output >= B_output/r.

Even a hypothetical channel transmitting one bit per Planck time would have rate

r approximately 1/t_P approximately 1.85 x 10^43 bits per second.

Transmitting 10^246 bits through such a channel would require approximately

T_output approximately 10^203 seconds.

This vastly exceeds the present age of the Universe.

One could introduce many parallel channels, but this only converts the time requirement into hardware and spatial requirements. To output all bits in one Planck-time interval would require approximately 10^246 simultaneously distinguishable output channels.

Parallelism changes the arrangement of the cost. It does not erase it.

6.3 Why Quantum Compression Cannot Print an Uncompressed Universe

Compression works when a long object can be generated from a shorter description.

A simple sequence such as

0101010101...

may be specified by a rule much shorter than the full string. A physical history governed by compact laws may likewise possess a description far smaller than its explicit spacetime raster.

Quantum states can also provide implicit representations. A wavefunction may encode a probability distribution over many configurations. A tensor network may represent a large state through local tensors and limited entanglement structure. A quantum circuit may generate a complex state from a compact sequence of gates.

None of these representations guarantees that every microscopic detail can be exported cheaply.

Consider a compressed file containing a trillion-bit image. The file may be much smaller than the expanded image, but viewing or printing every pixel requires decompression. The final expanded object still contains a trillion visible values.

The same distinction applies to a generative spacetime model.

A compact specification may take the form

D = initial state + laws + boundary conditions.

The explicit history may be produced by a reconstruction map

H = F(D).

The description length of D may be small relative to H. Yet evaluating F at every spacetime cell and recording every result may require work proportional to the size of H.

Description complexity and output complexity therefore differ.

A quantum computer might exploit D to answer a restricted question without constructing H in full. For example, it might estimate

<psi|O|psi>,

where O is an observable associated with a selected region or physical property.

The result might be one real number with an uncertainty interval. In that case, the requested output is small, and exhaustive reconstruction is unnecessary.

But if the requested observable is effectively the list of all local field values, then O has been replaced by an enormous family of measurements. The output advantage disappears.

Quantum compression is valuable only when the observer accepts a compressed answer, a global statistic, a sample, a decision, or a bounded collection of observables.

It cannot turn an uncompressed request into a compressed response while preserving every independent detail.

6.4 Measurement as an Irreversible Bottleneck

Measurement converts selected features of a quantum state into classical records. This transition imposes a bottleneck because a quantum system can support coherent combinations of alternatives while a classical record must contain definite, reproducible distinctions.

Suppose a system is prepared in

$|\psi\rangle = \sum_i a_i |i\rangle$.

A measurement in the basis $|i\rangle$ yields one outcome i with probability

$p_i = |a_i|^2$.

The observer does not receive the entire amplitude list. Repeated measurements provide samples. Different measurement bases can reveal different features of the state, but no single measurement extracts all noncommuting information.

The act of measurement also requires a physical apparatus. The outcome must be amplified, stabilized, and recorded. A microscopic distinction becomes correlated with a macroscopic state of a detector or memory.

This process creates a durable record only by coupling the quantum system to additional degrees of freedom. The readout device must itself possess enough distinguishable states to represent the possible outcomes.

If M distinct classical outcomes are required, the readout system must support at least M reliable records. The information capacity of the record is

$I_{record} = \log_2 M$.

For a complete N-bit spacetime image,

$M = 2^N$,

and therefore

$I_{record} = N$.

The readout apparatus must be at least as informationally expressive as the output it is expected to hold.

Measurement also introduces statistical cost. To estimate a probability or expectation value to finite precision, the experiment usually must be repeated. If the standard error decreases as

$1/\sqrt{M_{samples}}$,

then improving precision by a factor of ten requires approximately one hundred times as many samples.

Selected quantum estimation methods can improve this scaling for particular problems, but no method yields arbitrary precision from no resources.

The readout bottleneck therefore includes:

the number of distinct outputs,

the number of measurements,

the required precision,

the stability of the record,

and the physical capacity of the output medium.

A universal wavefunction may be mathematically rich, but a finite observer receives only finite records produced through finite interactions.

6.5 The No-Free-Readout Principle

The preceding arguments can be summarized in the No-Free-Readout Principle:

A quantum system may represent or process information implicitly within a high-dimensional state space, but extracting N independently specified classical bits requires an output process and record capable of supporting N classical bits.

The principle does not claim that every quantum computation requires work proportional to the dimension of the Hilbert space. That would deny the possibility of quantum advantage.

Instead, it draws a boundary between two kinds of tasks.

The first kind asks for a compact property of a large state:

a probability,

a period,

a correlation,

an energy eigenvalue,

a yes-or-no decision,

a sample,

or an expectation value.

Such tasks may sometimes be solved without enumerating every component of the state.

The second kind asks for an exhaustive classical description:

every amplitude,

every local value,

every branch,

every cell,

or every configuration weight.

These tasks contain large outputs by definition. Their cost cannot be removed merely by choosing a compact internal representation.

The No-Free-Readout Principle may be written schematically as

classical information extracted <= classical information physically recorded.

A related form is

full output cost >= output length.

This lower bound is elementary but powerful. It prevents a recurring misconception that quantum superposition, continuous amplitudes, or infinite-dimensional Hilbert spaces can release unlimited classical information through a finite measurement interface.

Applied to the spacetime raster, the principle gives a direct conclusion.

A quantum computer might represent a law-governed spacetime history implicitly.

It might simulate selected regions.

It might estimate global observables.

It might reconstruct local details on demand.

It might exploit correlations, symmetry, holography, and causal structure.

It cannot print a complete 10^246-cell classical movie without producing a correspondingly enormous record.

This conclusion changes the proper objective of quantum spacetime computation. The goal should not be exhaustive readout. It should be selective interrogation.

A physically meaningful machine would be asked questions such as:

What is the probability of a specified event?

What is the expected curvature in this region?

What correlations connect these two boundaries?

What macroscopic history is most consistent with these observations?

What local structure must be reconstructed to answer a defined measurement question?

These outputs may be finite even when the underlying state space is immense.

The productive use of Hilbert-space richness is therefore not to extract everything. It is to arrange interference, dynamics, and measurement so that the small portion extracted is the portion that matters.

The No-Free-Readout Principle does not diminish quantum computation. It identifies its proper domain. Quantum systems are not infinite encyclopedias waiting to be opened. They are physical processors capable of transforming hidden structure into selected observable answers.

7. Generative Spacetime Instead of Stored Spacetime

7.1 States, Laws, and Histories

A complete physical history need not be represented as an independent value assigned to every point in space and time. A more economical description begins with a state, a set of dynamical laws, and the conditions under which those laws operate.

In classical mechanics, an initial position and momentum may determine a later trajectory once the governing equations are specified. In classical field theory, field values and their rates of change on an appropriate spatial surface can determine subsequent evolution. In quantum mechanics, an initial state evolves under a Hamiltonian or a more general dynamical rule. In general relativity, the geometry of spacetime and the distribution of matter are linked through constraint and evolution equations.

These theories differ profoundly, but they share a generative structure:

present or boundary data + lawful evolution -> physical history.

The complete history is therefore not necessarily an independently stored sequence. It may be the result of applying a compact set of transformation rules to a suitable starting condition.

For a closed quantum system, the evolution may be written schematically as

$|\psi(t)\rangle = U(t, t_0)|\psi(t_0)\rangle,$

where |psi(t_0)> is an initial state and U is the evolution operator.

If the Hamiltonian H is time-independent, then

U(t, t_0) = exp[-iH(t - t_0)/hbar].

This expression does not list the state at every intermediate moment. It gives a rule from which the state at a selected time may be generated.

The distinction resembles the difference between a trajectory and the equation producing it. A trajectory sampled at 10^12 moments contains 10^12 entries. The governing equation and initial conditions may require far fewer symbols.

This does not imply that every physical history is easily compressible. Initial conditions may be complex. Quantum states may contain extensive entanglement. Chaotic dynamics may amplify tiny uncertainties. A compact law may generate an output whose detailed reconstruction requires enormous computation.

The central point is narrower: the number of spacetime samples is not automatically equal to the number of independent inputs needed to specify the history.

A generative representation of spacetime therefore separates four components:

the state from which evolution begins,

the laws governing transformation,

the boundary and consistency constraints,

and the observables selected for reconstruction.

The resulting history may be vast, while its physical specification is structured.

7.2 The Formal Generative Description

The raster model represents spacetime as a collection of local entries:

R = {s(x, y, z, t)}

for every allowed coordinate tuple.

If there are approximately 10^246 spacetime cells and each cell contains b bits, then the explicit raster has size

B_raster approximately b x 10^246.

A generative model instead represents spacetime through a tuple such as

G = (psi_0, L, C, B, Q),

where

psi_0 = initial or boundary state,

L = dynamical laws,

C = physical constraints,

B = boundary conditions,

Q = requested observables or queries.

The output associated with a query q may then be written as

A_q = F(G, q),

where F is a reconstruction, simulation, inference, or measurement procedure.

This representation changes the computational objective. The machine is no longer asked to produce the universal raster R. It is asked to calculate A_q for a defined question q.

The query may specify:

a region of spacetime,

a resolution scale,

a physical observable,

a time interval,

a confidence requirement,

and an acceptable approximation error.

For example, the machine might be asked to estimate the expected energy density within a defined region, rather than report the value of every microscopic field variable inside it.

The generative description is especially useful when the output contains strong regularity. A homogeneous region need not be represented by a separate independent value at every point. A symmetry rule plus a small set of parameters may suffice.

Similarly, a propagating wave need not be stored as unrelated samples at every location and time. Its amplitude, frequency, phase, direction, dispersion law, and interactions may provide a compact description.

The compression can be expressed schematically as

description length << raster length

when the history contains sufficient lawful structure.

However, the existence of a compact description does not guarantee inexpensive reconstruction. A short rule may generate behavior that is computationally irreducible, meaning that obtaining the detailed state at a late time may require effectively following the evolution step by step.

The generative model therefore separates three quantities:

K_description = size of the state, laws, and constraints,

C_reconstruction = computation required to answer a query,

B_output = size of the requested classical result.

These quantities need not scale together.

A history may have a short description but require enormous reconstruction time. A complex state may nevertheless permit efficient calculation of a simple global observable. A small output may require less computation than a full local rendering.

The tuple G should therefore not be mistaken for a magical compression that makes all questions easy. It is a framework for asking which parts of the history must actually be generated.

7.3 Temporal Redundancy

The multiplication

10^185 spatial cells x 10^61 temporal intervals approximately equals 10^246 spacetime cells

implicitly treats every temporal slice as an independent data object. Physical evolution does not generally have this structure.

If the state at time t + Delta t is determined by the state at time t and a dynamical rule, then much of the information in neighboring slices is redundant.

A discrete evolution may be written as

S_(n+1) = F(S_n),

where S_n is the state at temporal step n.

The full sequence is

S_0, S_1, S_2, ..., S_N.

An explicit archive stores all N + 1 states. A generative description may store only

$S_0 + F + N$.

The sequence can then be reconstructed as needed.

In a reversible system, earlier states may also be recoverable from later ones:

$S_n = F^{-1}(S_{n+1})$.

Quantum evolution of a closed system is unitary and therefore reversible at the level of the complete state. This suggests that successive universal quantum states are not independent additions of new arbitrary information. They are related by structured transformations.

The apparent information growth in macroscopic history often reflects coarse-graining, decoherence, record formation, and the practical loss of access to microscopic correlations. The underlying closed-system evolution may remain reversible even when thermodynamic behavior appears irreversible.

Temporal redundancy can be dramatic when change is slow. Most local variables may differ only slightly between adjacent Planck-time intervals. Recording the full value repeatedly is analogous to storing every frame of a motionless scene without compression.

A differential description stores rates of change rather than complete replacement states. For a field ϕ,

$$\frac{\partial \phi}{\partial t} = F(\phi, \text{spatial derivatives}, \text{sources}).$$

Given appropriate initial data, the equation specifies how later values arise.

Yet temporal compression has limits. Sensitive dependence on initial conditions can make long-term prediction difficult. Quantum measurements may introduce branching descriptions relative to observers. Open systems exchange information with their environments. Singularities, phase transitions, or topology changes may challenge simple evolution rules.

Even so, none of these complications justifies assuming that every time slice contains a wholly independent set of arbitrary values.

The Planck-time count is therefore better interpreted as the maximum number of temporal sampling opportunities than as the number of independent universal states that must be stored simultaneously.

A generative spacetime computer would exploit this redundancy. It might retain:

the initial state,

periodic checkpoints,

macroscopic records,

error-correction data,

and selected observables.

It would not necessarily preserve every microscopic intermediate state as a separate classical frame.

7.4 Gauge Redundancy and Equivalent Descriptions

Physical theories often contain mathematical variables that do not correspond one-to-one with observable distinctions. Gauge redundancy means that several mathematical descriptions may represent the same physical state.

In electromagnetism, different scalar and vector potentials can produce the same electric and magnetic fields. The potentials are useful for calculation, but not every difference between potential functions corresponds to a physically distinct situation.

In general relativity, coordinate labels do not possess independent physical meaning. The same spacetime geometry may be expressed using many coordinate systems. A change of coordinates can alter the numerical form of the metric components without changing the underlying physical geometry.

A naive spacetime raster may count such representational differences as if they were independent information.

Suppose every cell stores coordinate-dependent field components. The raw bit count includes:

physical degrees of freedom,

coordinate choices,

gauge choices,

constraint-related variables,

and numerical redundancy.

The physically meaningful information is obtained only after equivalent descriptions are identified.

This may be expressed as

physical state space = mathematical configuration space / gauge equivalence.

The division symbol here indicates that configurations connected by gauge transformations are treated as physically equivalent.

The reduction can be substantial. A mathematical field may have many components, while constraints and gauge symmetry leave only a smaller number of propagating physical degrees of freedom.

In electromagnetism, the four-component potential does not imply four independent physical polarizations of a photon. Gauge freedom and constraints reduce the observable propagating modes to two transverse polarizations.

In general relativity, the spacetime metric has ten symmetric components in four dimensions, but coordinate freedom and constraint equations greatly reduce the number of independent propagating gravitational degrees of freedom.

A Planck-scale database that stores all metric components independently at every cell would therefore overcount physical information. Many entries would be related through constraints or removable through coordinate choice.

Gauge redundancy is more than a technical correction to the voxel estimate. It challenges the idea that local coordinate values are the fundamental contents of spacetime.

If points are identified only through relational or invariant properties, then an isolated coordinate address such as

(x, y, z, t)

may not represent a physically meaningful cell without reference to matter, geometry, observers, or boundary conditions.

A generative spacetime description should therefore prioritize gauge-invariant quantities, relational observables, causal relationships, and boundary data. It should avoid treating every coordinate-dependent variable as an independently stored fact.

The reduction from mathematical variables to physical observables is one of the mechanisms by which the apparent information content of a spacetime raster may greatly exceed the number of true degrees of freedom.

7.5 Causality and Local Relevance

Relativity imposes a powerful restriction on what information can affect a given event. An observation at event P can depend only on conditions within its causal past, subject to the propagation rules of the relevant fields.

Events outside the past light cone of P cannot influence P through ordinary causal processes.

This permits a major reduction in query-specific computation. To answer a question about a local event, a simulator need not reconstruct the entire Universe at Planck resolution. It needs only the portion of the state that can causally contribute to the requested observable.

Let C^-(P) denote the causal past of P. Then a local prediction may be written schematically as

A(P) = F[data in C^-(P), laws, boundary conditions].

Information outside C^-(P) can be excluded unless it contributes through global constraints, entanglement structure, or boundary conditions relevant to the problem.

For an extended spacetime region Omega, the relevant computational domain becomes its causal past:

C^-(Omega).

This causal restriction suggests an adaptive architecture. A quantum spacetime compiler would begin with the requested observable and trace backward through the causal network to identify relevant degrees of freedom.

It might then:

discard spacelike-separated regions that cannot affect the answer,

coarse-grain distant sources whose fine structure is irrelevant,

retain high resolution near sensitive interactions,

and evolve only the necessary causal domain.

This is analogous to selective rendering in computer graphics. A graphics engine does not calculate every possible viewpoint at maximum resolution. It renders the scene visible to the camera. A spacetime compiler would render only the physics visible to the selected observable.

Quantum entanglement complicates the picture because correlations can extend across spacelike separation. Entanglement does not permit controllable faster-than-light signalling, but joint observables may depend on global quantum structure. A valid reduction must therefore preserve the correlations required for the question even when no direct causal signal crosses between the regions during the measurement interval.

The relevant domain is not always a simple classical light cone. It may include entangled subsystems, constraint surfaces, boundary conditions, and encoded nonlocal information.

Nevertheless, the principle of local relevance remains strong:

not every degree of freedom is relevant to every question.

The computational cost should therefore be associated with the causal and informational support of the observable, not automatically with the total number of cells in the Universe.

This leads to a more useful formulation of the spacetime-computation problem.

The exhaustive raster asks:

What is the value of every variable everywhere and everywhen?

The generative approach asks:

What subset of states, laws, correlations, and causal pathways must be represented to answer this particular physical question to the requested accuracy?

The first demand is universally maximal. The second may be finite and tractable for selected cases.

Generative spacetime therefore replaces universal storage with conditional reconstruction. The full physical history remains implicit in the combination of state, dynamics, constraints, and correlations. A finite observer accesses only a bounded projection of that history.

The concept may be summarized as

observable reality for a finite query =
relevant state

- governing dynamics
- causal domain
- measurement rule.

This does not prove that spacetime itself is literally computed on demand. It establishes a more economical and physically grounded way to represent spacetime within a theory or simulator.

The Universe need not be treated as a movie containing every frame in explicit storage. It may be represented as a constrained quantum process from which particular histories, correlations, and observations are generated.

Once this perspective is adopted, the fundamental challenge is no longer to build a memory large enough for 10^246 entries. It is to identify the smallest generative structure sufficient to reproduce every physically meaningful answer available to an observer.

8. Holography, Entropy Bounds, and Spacetime Compression

8.1 Volume Counting Versus Boundary Counting

The Planck raster begins with volume. A three-dimensional region is divided into cubes of side l_P, and the number of cells is estimated by

N_volume approximately V/l_P^3.

For the observable Universe, this produces a spatial count on the order of 10^185 Planck-volume cells. Extending the construction across approximately 10^61 Planck-time intervals produces a four-dimensional count near 10^246.

This procedure assumes that the amount of independent information in a region scales with the number of small volume elements inside it. That assumption is natural in ordinary lattice models. If every lattice site can vary independently, doubling the volume approximately doubles the number of local degrees of freedom.

Gravity challenges this intuition.

Black-hole thermodynamics and related entropy bounds indicate that the maximum physically meaningful entropy in a bounded region may scale with the area of its boundary rather than with its volume. In Planck units, the relevant count has the general form

N_area approximately A/l_P^2,

where A is the boundary area.

For a sphere of radius R,

A = 4pi R^2.

Using a cosmic scale of approximately

R approximately 4.4 x 10^26 meters,

the corresponding area is on the order of

A approximately 2.4 x 10^54 square meters.

The Planck area is

l_P^2 approximately 2.61 x 10^-70 square meters.

The ratio is therefore approximately

A/l_P^2 approximately 9 x 10^123.

Factors such as 1/4, the choice of horizon, and the conversion between thermodynamic entropy and bits reduce the characteristic information estimate to the familiar broad scale near

10^{122} to 10^{123} bits.

The precise coefficient is not the main issue. The important difference is the scaling law:

volume count proportional to R^3,

area count proportional to R^2.

For a very large region, the gap between these quantities becomes enormous.

If the maximum independent information associated with a gravitational region scales with its boundary area, then a count of Planck-volume cells grossly overestimates the number of independent physical degrees of freedom. The interior may contain many local descriptions, but they cannot all vary independently while remaining consistent with gravity.

This does not mean that the interior is empty or unreal. It means that the information required to distinguish physically allowed states may be much smaller than the number of points used to describe the interior geometry.

A smooth curve can be sampled at millions of points while being generated by a small equation. Similarly, a gravitational region may admit an enormous number of local coordinate values while possessing a more economical physical description constrained by boundaries, symmetries, and dynamics.

The shift from volume counting to boundary counting is therefore not merely a numerical correction. It is a change in the assumed architecture of physical information.

8.2 Black-Hole Entropy as an Information Limit

A black hole provides the strongest known indication that information in a finite region cannot increase without bound.

Suppose energy is added to a bounded region in an attempt to store more information. Additional memory states require physical distinctions, and physical distinctions require matter, energy, field configurations, or other resources. If enough energy is concentrated within a sufficiently small region, gravitational collapse occurs.

The resulting black hole has an entropy proportional to the area of its event horizon:

$S_{BH} = k_B A/(4l_P^2)$,

where k_B is Boltzmann's constant.

The corresponding number of binary information units is approximately

$I_{BH} = S_{BH}/(k_B \ln 2)$

and therefore

$I_{BH} = A/(4l_P^2 \ln 2)$.

The striking feature is that the entropy is determined by area, not interior volume.

A black hole is widely regarded as the most entropic object that can occupy a region of a given boundary area. If a proposed memory configuration inside that region had more entropy than a black hole of the same size, gravitational collapse would appear to produce an object with less entropy than the original system. That would conflict with the generalized second law unless the assumed interior state count was physically redundant or inaccessible.

The implication is that one cannot indefinitely increase usable information density by packing more microscopic states into a fixed volume. Gravity eventually changes the geometry and places a horizon around the system.

The horizon limits what an external observer can distinguish. Many possible interior descriptions may correspond to the same externally measurable mass, charge, angular momentum, horizon area, and radiation pattern. The external information capacity is therefore not simply the count of imagined interior lattice sites.

This is directly relevant to the proposed Planck-scale computer. A machine containing one independent degree of freedom for every Planck-volume cell would require enormous energy and physical structure. If enough independent information were concentrated into a finite region, the computer would itself become a gravitating object and ultimately a black hole.

At that point, the computer could no longer be treated as a passive information container embedded in an unchanged spacetime. Its energy would alter the geometry it was intended to represent.

The limit is therefore self-enforcing:

information requires physical states,

physical states require energy and structure,

energy and structure gravitate,

and excessive concentration produces a horizon.

The machine cannot indefinitely add memory without becoming part of the gravitational problem.

8.3 Cosmological Information Scales

Applying black-hole and horizon reasoning to the observable Universe is more subtle than applying it to an isolated black hole. Cosmology contains several possible horizons, including particle horizons, apparent horizons, Hubble scales, and event horizons. Their meanings differ, and not every cosmological spacetime has the same global causal structure.

Nevertheless, a characteristic horizon area in Planck units often produces an information scale near

10^122 bits,

within one or more orders of magnitude depending on the chosen definition and assumptions.

This number is much smaller than the approximately 10^185 Planck-volume cells in one spatial raster and vastly smaller than the approximately 10^246 cells in a complete four-dimensional raster.

The comparison is approximately

10^185/10^122 = 10^63

for one spatial slice,

and

10^246/10^122 = 10^124

for the full spacetime raster.

These ratios suggest that the naive voxel representation may contain between 10^63 and 10^124 times more entries than the number of independent bits permitted by a horizon-scale gravitational bound.

The interpretation must be handled carefully.

The 10^122 figure is not an experimentally counted memory register. It is an order-of-magnitude gravitational entropy scale. It does not mean that every cosmological degree of freedom has been fully identified or that the entire Universe can be represented by a conventional file containing exactly 10^122 bits.

Nor does it establish that the observable Universe is a black hole. Cosmological horizons and black-hole horizons arise in different geometries and require different treatments.

The comparison remains important because it shows that gravity does not support arbitrary volume-based information counting. The Planck raster may be a useful coordinate sampling, but it cannot automatically be interpreted as a set of independent physical memory cells.

The discrepancy also suggests that time should not be counted as an independent storage axis in the same way as space. A four-dimensional history generated by lawful evolution does not necessarily contain one new independent set of bits at every time step.

If a cosmic state has an information capacity of order 10^122 bits at a given time, multiplying it by 10^61 time samples does not necessarily produce 10^183 or 10^246 independent bits. The later states are related to earlier states by dynamics, constraints, and causality.

The four-dimensional raster is therefore best understood as a sampling burden, not an entropy count.

8.4 The Nominal Compression Gap

The difference between the raster estimate and the horizon-scale estimate may be called the nominal compression gap.

Let

N_raster approximately 10^246

and

I_horizon approximately 10^122 bits.

Then

C_gap = N_raster/I_horizon

gives

C_gap approximately 10^124 cells per horizon-scale bit.

This ratio should not be interpreted literally as a practical compression algorithm. It does not mean that each block of 10^124 spacetime cells can be replaced by one ordinary bit through a known encoding.

Rather, it quantifies the scale of the mismatch between two models:

the raster model, in which each Planck four-cell is treated as independently recordable,

and the gravitational model, in which the number of independent states is limited by a boundary area.

The gap indicates that most raster entries cannot be independent physical variables.

They may be constrained through several mechanisms:

successive time samples may be dynamically related,

neighboring spatial values may be correlated,

gauge-equivalent configurations may describe the same physics,

local fields may be determined by boundary conditions,

high-frequency modes may be physically inaccessible,

entanglement may organize information nonlocally,

and gravitational constraints may exclude arbitrary combinations.

The compression gap can also be viewed as a challenge to ontology. If a coordinate raster contains 10^246 cells but the physical state capacity is closer to 10^122 bits, then most of the raster cannot correspond to separate elementary facts.

The cells may instead be reconstruction points within an emergent geometry.

A map may contain billions of pixels without the terrain containing one independent object for each pixel. The pixels are features of the representation. Similarly, Planck-scale cells may be features of a hypothetical discretization rather than fundamental storage locations.

This suggests a hierarchy:

coordinate complexity,

field-description complexity,

physical-state complexity,

observable complexity.

The raster count belongs to the first or second level. The entropy bound concerns the third. A finite measurement addresses the fourth.

Confusing these levels produces the illusion that the Universe must physically store every coordinate sample as a separate datum.

8.5 Description Complexity Versus Reconstruction Complexity

A compressed description can be much shorter than the object it generates.

Consider a sequence of one trillion zeros. Its explicit length is one trillion bits, but its description may be written compactly as

print 0 one trillion times.

The description is short. The output remains long.

This distinction becomes more important when the generated object is complex. A short computer program can produce a sequence whose individual entries are difficult to predict without executing the program step by step. The rule is compact, but reconstruction remains computationally expensive.

Spacetime may exhibit the same separation.

Let

K_history

denote the shortest description of the laws, boundary data, and initial state sufficient to generate a history.

Let

C_history

denote the computational work required to reconstruct that history at a chosen resolution.

Let

B_history

denote the size of the explicit output.

These quantities may satisfy

K_history << C_history <= B_history,

although their exact relationship depends on the problem.

A holographically compressed or boundary-based representation could greatly reduce K_history without making C_history small.

For example, a boundary state may implicitly encode a complicated interior geometry. Recovering a particular interior observable may require a difficult decoding procedure. Recovering every local feature could remain effectively impossible.

This distinction prevents an overly optimistic interpretation of holography. A lower information bound does not automatically imply an easy simulation.

Compression answers the question

How little information is sufficient to specify the allowed state?

It does not necessarily answer

How quickly can every local detail be reconstructed?

Nor does it answer

How much classical output is required to display the reconstruction?

A state may be compactly specified yet computationally irreducible. The only way to know its detailed late-time behavior may be to evolve it through many intermediate stages.

A quantum spacetime compiler would therefore need to exploit query structure. It should reconstruct only those details needed for the requested observable.

If the query asks for a horizon area, a global energy, or a long-distance correlation, a coarse or boundary description may suffice.

If the query asks for every local field value at every Planck time, the reconstruction cost may return toward the original raster scale.

The value of compression lies in making selected questions feasible, not in making exhaustive reality free.

8.6 The Spacetime Compression Hypothesis

The preceding arguments motivate the Spacetime Compression Hypothesis:

The physically independent information content of a bounded spacetime region is far smaller than the number of Planck-scale coordinate cells used to sample that region, because lawful evolution, gauge equivalence, causality, entanglement, and gravitational entropy bounds impose extensive redundancy and constraint.

The hypothesis does not require spacetime to be literally digital. It applies whether the underlying theory is continuous, discrete, algebraic, causal, holographic, or emergent.

Its core claim is that physical reality should not be equated with an exhaustive list of local coordinate values.

A compressed spacetime representation may instead take the schematic form

spacetime description =
boundary information

- dynamical laws
- quantum state
- consistency constraints
- reconstruction map.

The reconstruction map would generate local geometry and field behavior only when a particular observable or region is queried.

In this view, the interior is not absent. It is encoded relationally and dynamically rather than stored as an independent classical file.

The hypothesis has several consequences.

First, the maximum useful information capacity of a region should be estimated through physical distinguishability and gravitational bounds, not merely through lattice-site counting.

Second, time should be treated as lawful transformation rather than as a sequence of independently stored frames.

Third, local geometry may be emergent from more fundamental quantum relationships, making the Planck voxel a derived concept rather than a primitive memory element.

Fourth, a simulator should organize information according to causal relevance, boundary structure, and observables rather than uniform four-dimensional resolution.

Fifth, the correct measure of computational difficulty may depend more on entanglement structure, decoding complexity, and query precision than on raw coordinate volume.

The Spacetime Compression Hypothesis does not solve quantum gravity. It reframes the computational question in a way consistent with the strongest clues supplied by gravitational thermodynamics.

The original Planck raster asks how many cells would be required to print reality.

Holographic reasoning asks how many independent states reality can physically possess.

These are not the same question.

The gap between them suggests that the Universe is not fundamentally organized as a four-dimensional memory array. It is more plausibly represented as a constrained quantum system in which geometry, matter, and information are encoded through relationships whose independent capacity scales far more slowly than a naive spacetime-cell count.

The practical implication is decisive. A viable quantum spacetime computer should not attempt to allocate one register to every Planck cell. It should seek the compressed variables from which physically meaningful interior observables can be reconstructed.

The next step is therefore not to build a larger raster machine. It is to design a quantum spacetime compiler capable of translating boundary data, causal structure, dynamical laws, and selected queries into finite observable answers.

9. The Quantum Spacetime Compiler

9.1 From Universal Rendering to Bounded Questions

A literal Planck-scale simulator would attempt to calculate and display every local variable throughout the observable Universe at every Planck-time interval. The previous sections showed why this objective is physically and computationally misguided. Its output alone would contain approximately 10^246 spacetime cells before any realistic description of fields, geometry, uncertainty, or quantum correlations was added.

A quantum spacetime compiler would pursue a different objective. It would not construct a complete microscopic movie of reality. It would translate a bounded physical question into the smallest quantum computation capable of answering that question to a specified accuracy.

The difference resembles the distinction between rendering an entire virtual world at maximum resolution and rendering only the portion visible to a particular observer. A graphics engine does not calculate every possible viewpoint, every hidden surface, and every wavelength of light. It identifies the scene, camera, illumination, and resolution relevant to the requested image.

A spacetime compiler would perform an analogous operation for physical observables.

Instead of asking

What is happening at every spacetime cell?

it would ask

What information must be represented to determine this observable?

The query might concern the probability of a particle arriving at a detector, the correlation between two regions, the curvature along a worldline, the likelihood of black-hole formation, or the range of initial conditions compatible with a measured final state.

Such questions can have finite outputs even when the underlying state space is enormous.

The compiler would therefore transform a universal problem into an observable-specific problem:

universal state + bounded query -> finite answer.

This does not make every quantum-gravity problem tractable. Some bounded questions may still require enormous resources. The point is that the computation should scale with the information relevant to the question rather than automatically with the full Planck raster.

9.2 Inputs to the Compiler

A quantum spacetime compiler would require a carefully defined input specification. At minimum, the input would contain an initial or boundary quantum state, a set of dynamical laws, a spacetime domain, an observable, a resolution scale, and an accepted error tolerance.

The input may be represented schematically as

I = (psi, L, Omega, O, r, epsilon),

where

psi = initial or boundary state,

L = dynamical laws and constraints,

Omega = causal or geometric domain,

O = requested observable,

r = required spatial, temporal, or spectral resolution,

epsilon = tolerated uncertainty or approximation error.

The initial state psi need not be a complete classical configuration. It might be a wavefunctional, density operator, boundary state, tensor-network description, effective field configuration, or probabilistic ensemble.

The laws L would specify the permitted transformations. Depending on the theory, they could include a Hamiltonian, action, constraint algebra, causal update rule, scattering operator, or effective dynamics.

The domain Omega would identify the physical region relevant to the problem. It could be a laboratory volume, a causal diamond, a cosmological horizon, a black-hole exterior, or a finite interval surrounding an interaction.

The observable O would define what the machine must return. This is essential because quantum theory does not supply all possible properties as simultaneously readable values. A

calculation must specify the measurement operator, correlation function, transition amplitude, or probability distribution of interest.

The resolution r would prevent unnecessary refinement. A question about galactic structure does not ordinarily require Planck-scale detail. A query involving high-energy scattering may require finer resolution locally but not across the entire Universe.

The error tolerance epsilon would establish when the answer is sufficient. Exact calculation is generally impossible for large physical systems. A useful compiler must quantify numerical error, model uncertainty, truncation error, and measurement uncertainty.

The resource demand may then be expressed as

C_query = C(psi, L, Omega, O, r, epsilon).

This formulation replaces the assumption of a single universal computational cost with a family of query-dependent costs.

9.3 Causal-Cone Reduction

The first major operation of the compiler would be causal reduction.

An observable associated with an event or region can be influenced only by information capable of reaching that event or region according to the governing causal structure. The compiler should therefore identify the relevant causal past rather than evolve the entire Universe.

For an observable supported in region Omega, let

C^-(Omega)

denote its causal past. The calculation would initially be restricted to information within this domain.

This restriction can reduce the problem dramatically. A detector operating for one second in a laboratory does not require a detailed simulation of every galaxy. Distant matter may contribute through aggregate gravitational or electromagnetic effects, but its microscopic state can often be replaced by effective boundary conditions or coarse multipole descriptions.

The compiler would trace dependency pathways backward from the observable. It would ask which interactions, sources, fields, and boundary data can alter the requested result within the accepted error tolerance.

Degrees of freedom would then be classified as:

directly relevant,

indirectly relevant through interactions,

relevant only through effective parameters,

or negligible at the requested precision.

The resulting active domain may be much smaller than the original spacetime region.

Quantum entanglement requires special treatment. Spacelike-separated systems can exhibit correlations even though they cannot transmit controllable signals faster than light. A causal reduction must preserve entangled partners when their correlations affect the selected observable.

The relevant computational domain is therefore not always identical to a classical geometric light cone. It may include nonlocal encoded correlations, gauge constraints, boundary relationships, and global topological sectors.

A more complete relevance region may be written as

R(O) = causal support + correlation support + constraint support.

The compiler would retain all three while discarding variables that cannot materially affect the answer.

9.4 Adaptive Resolution and Multiscale Representation

Uniform Planck-scale resolution is one of the main sources of the impossible 10^246-cell estimate. Most physical questions do not require the same resolution everywhere.

A quantum spacetime compiler would use adaptive resolution.

Regions containing strong curvature, rapid field variation, high-energy interactions, singular behavior, or sensitive interference could be represented finely. Regions with smooth fields or weak influence could be coarse-grained.

The resolution would become a function of location and scale:

r = r(x, t).

A black-hole horizon might require high geometric resolution, while a distant homogeneous region might be represented by a few macroscopic parameters.

Adaptive refinement would proceed iteratively. The compiler would begin with a coarse description and estimate whether unresolved structure could change the observable by more than epsilon. If so, the relevant region would be refined.

The process could be written schematically as

coarse model -> error estimate -> selective refinement -> updated answer.

This resembles adaptive mesh refinement in classical simulation, but the quantum version would also need to manage entanglement, state truncation, basis selection, and error correction.

Multiscale representation would allow the same physical system to be described differently at different scales. Microscopic particle or field variables might be used near an interaction, while hydrodynamic, thermodynamic, or geometric variables describe distant regions.

Renormalization provides the conceptual basis for such scale changes. Short-distance degrees of freedom can often be absorbed into effective parameters governing long-distance behavior.

The compiler should therefore avoid resolving microscopic detail unless that detail survives into the observable.

A query about the gravitational field of a star does not ordinarily require the exact quantum state of every atom in the star. Its mass, angular momentum, multipole moments, and selected fluctuations may suffice.

The computational advantage comes from matching representation to relevance.

9.5 Symmetry, Conservation, and Constraint Exploitation

Physical law contains extensive structure that can reduce computation.

Symmetry identifies states or configurations related by transformations. If a system is rotationally symmetric, many angular configurations are equivalent. If it is translationally invariant, the same local rule applies at many positions.

A compiler should detect and exploit such structure rather than calculate equivalent cases independently.

Conservation laws further restrict possible histories. Energy, momentum, charge, angular momentum, and other conserved quantities define sectors of the state space. If the initial state occupies one sector, the computation need not explore sectors that cannot be reached.

Gauge constraints remove representational redundancy. Coordinate choices, phase conventions, and nonphysical field components should not consume independent memory when they describe the same physical state.

Locality also permits decomposition. A local Hamiltonian acts through interactions among nearby degrees of freedom. This allows the compiler to use sparse operators, tensor decompositions, causal circuits, and domain-specific factorizations.

The active state space can therefore be reduced through a sequence such as

formal configurations

-> constraint-satisfying configurations

-> gauge-inequivalent configurations

-> causally relevant configurations

-> observable-sensitive configurations.

Each stage removes distinctions that cannot affect the requested answer.

Conservation and symmetry may also enable exact calculations for quantities that would otherwise require full simulation. A global charge can sometimes be inferred from boundary flux. A transition may be forbidden by symmetry without calculating every microscopic path. A conserved energy may remain fixed even while the local state becomes highly complex.

The compiler should prefer such structural inference whenever possible.

9.6 Observable-Oriented Output

The final output of the compiler would be tailored to the query rather than presented as an exhaustive spacetime reconstruction.

Possible outputs include

a transition probability,

an expectation value,

a correlation function,

a scattering amplitude,

a causal-connectivity statement,

a range of compatible histories,

an estimated horizon area,

or a coarse-grained geometric configuration.

A result might take the form

P(event X | conditions C) = 0.731 +/- 0.004.

Another query might return

= K +/- Delta K.

A correlation query could produce

C(A, B) = <O_A O_B> - <O_A><O_B>.

These answers are finite summaries of selected features of a much larger quantum state.

The compiler should also report the assumptions and limits of the answer. A physically meaningful result would include the resolution scale, model truncations, error estimate, causal domain, and any dependence on uncertain boundary data.

The output should distinguish several forms of uncertainty:

statistical measurement uncertainty,

numerical approximation error,

state-preparation uncertainty,

model uncertainty,

and uncertainty arising from incomplete boundary conditions.

The machine would therefore produce not only a value but an epistemic description of that value.

In some cases, the correct output may be a sample rather than a full distribution. Quantum systems naturally generate outcomes according to probabilities. A collection of samples may be sufficient for testing a hypothesis, estimating a parameter, or comparing a theoretical model with observation.

In other cases, the output may be a compressed representation that can be queried repeatedly. A tensor network, reduced density matrix, effective action, or causal graph might summarize the relevant state without expanding it into local values everywhere.

The central design principle is that output should remain proportional to the scientific question.

A machine asked whether a black hole forms should not first print every Planck-scale field value in the Universe.

A machine asked for a local transition amplitude should not reconstruct the full cosmic history.

A machine asked for an observable should calculate the observable.

The quantum spacetime compiler is therefore not a universal camera. It is an interface between physical law and finite observation.

Its purpose is to transform an immense implicit quantum reality into bounded answers that can be prepared, computed, measured, and understood by finite observers.

The concept can be summarized as

quantum spacetime compiler =
query specification

- causal reduction
- adaptive representation
- structural compression
- observable extraction.

Such a machine would not overcome the No-Free-Readout Principle. It would respect it. Its advantage would come from refusing to request unnecessary output.

The success of the compiler would be measured not by the amount of spacetime it renders, but by the amount of irrelevant computation it avoids while preserving the accuracy of the requested physical answer.

10. Physical Limits on Hilbert-Space Harvesting

10.1 Finite Energy

Every usable quantum state must be embodied in a physical system, and every physical system operates within a finite energy budget. This immediately limits the portion of an infinite-dimensional Hilbert space that can be prepared, controlled, and measured.

For a harmonic oscillator, the energy levels are

E_n = hbar omega(n + 1/2).

Although n has no formal upper bound, access to higher values requires increasing energy. If the available energy is E_max, then the largest accessible occupation number is approximately

n_max approximately E_max/(hbar omega) - 1/2.

The operational dimension is therefore finite even when the formal Hilbert space is infinite.

One might lower omega to reduce the energy spacing between levels. This does not eliminate the cost. As adjacent levels become more closely spaced, they become harder to distinguish.

Longer measurement times, narrower spectral resolution, greater environmental isolation, and more accurate control are required.

Finite energy also limits spatial resolution. Probing shorter wavelengths requires larger momenta and therefore larger energies. A characteristic relation is

lambda approximately h/p.

As lambda decreases, the required momentum p increases. At sufficiently small scales, the energy used to probe the region contributes significantly to the gravitational field. The measuring apparatus can no longer be treated as external to the system being measured.

Energy constraints also restrict the number of particles, modes, and field excitations that can be maintained. A formal quantum field contains infinitely many modes, including arbitrarily high-frequency modes. A finite-energy state excites only a restricted portion of that spectrum to appreciable occupation.

This distinction can be written as

formal field space = infinitely many possible modes,

accessible field space = modes and occupations permitted by finite energy.

A quantum computer that attempted to exploit every formal mode would need unbounded excitation energy. A practical machine instead imposes cutoffs on frequency, occupation number, amplitude, and spatial resolution.

The energy requirement is not limited to the stored state. Quantum control requires pulses, fields, clocks, amplifiers, detectors, cooling systems, and classical processing. Error correction requires repeated measurements and active interventions. Communication among components requires signal energy.

Even logically reversible quantum gates must be implemented by finite physical interactions. Reversibility may reduce the minimum thermodynamic dissipation associated with certain logical operations, but it does not make control, stabilization, and readout costless.

The usable Hilbert harvest is therefore constrained by a total energy accounting:

E_total =
E_state

- E_control
- E_cooling
- E_measurement

- E_error correction
- E_communication.

A proposal that counts only the energy of the ideal quantum state while ignoring the supporting apparatus greatly overestimates accessible capacity.

For Planck-scale spacetime computation, finite energy is decisive. A register containing one independent qubit for every Planck-scale spatial cell would require approximately 10^185 logical degrees of freedom. Any nonzero average energy per degree of freedom would produce a total energy vastly beyond what can be assembled without gravitational consequences.

The problem cannot be solved merely by making each qubit individually weak. As energy gaps shrink, thermal excitation, noise, and slow control make the states less distinguishable. Information storage requires some physically resolvable separation between alternatives.

Finite energy therefore limits both how many states can be reached and how reliably they can be distinguished.

10.2 Finite Time and Quantum Speed Limits

A large Hilbert space is useful only if the desired regions of it can be reached within the available time. State preparation and computation require physical evolution, and physical evolution cannot proceed arbitrarily fast at finite energy.

Quantum speed limits relate the minimum time required for a system to evolve between distinguishable states to quantities such as its average energy above the ground state or its energy uncertainty.

A schematic form is

T_min greater than or approximately equal to hbar/Delta E,

where Delta E represents a relevant energy scale or uncertainty.

The exact bound depends on the initial state, target state, Hamiltonian, and definition of distinguishability. The general principle is that faster transformation requires greater energetic resources.

This creates an energy-time tradeoff. A device may attempt to gain access to more states by increasing its available energy, but doing so increases heating, control difficulty, and gravitational back-reaction. Reducing the energy lowers those costs but slows state preparation and measurement.

For a computation requiring G sequential logical steps, the runtime satisfies

T_run greater than or approximately equal to G tau_gate,

where tau_gate is the minimum reliable logical gate time.

If

G approximately 10^246,

even an unimaginably small gate time leaves an impossible sequential duration unless enormous parallelism is introduced.

Parallelism cannot remove causal depth. When later operations depend on earlier results, a minimum number of sequential layers remains. A computation may contain many gates that can be executed simultaneously, but dependency chains cannot be compressed below their logical depth.

The one-slice spacetime simulation described earlier requires approximately 10^61 temporal updates if every Planck-time interval is represented. If each simulated step depends on the preceding state, those updates form a sequential chain.

A perfect simulator evolving one Planck step per actual Planck time would merely keep pace with the process it simulates. It would not produce the complete cosmic history faster than the Universe itself.

Attempting to run faster raises a deeper issue. A simulator embedded within the Universe must itself obey the physical laws governing time, energy, and signal propagation. It cannot execute arbitrarily many distinguishable transformations within an arbitrarily short interval.

State preparation may impose an even greater cost than state evolution. A generic state in a Hilbert space of dimension d may require a circuit whose complexity grows rapidly with d. The fact that a state mathematically exists does not imply that a short sequence of available operations can produce it.

This motivates a distinction between

state-space size

and

state-reachability complexity.

A huge Hilbert space may contain states separated from the available initial state by circuits of extraordinary depth. Those states are operationally irrelevant if they cannot be reached before the apparatus decoheres, fails, or the available cosmological time expires.

The operational Hilbert dimension should therefore include a reachability condition:

d_op = number of states reachable and distinguishable within time T.

Time also limits readout. Estimating small probability differences requires repeated measurements. Resolving narrow spectral features requires long observation. Reconstructing a complicated state requires many preparations and measurement settings.

Finite time filters the formal state space just as strongly as finite energy.

10.3 Finite Size and Communication

A quantum computer is not an abstract point. Its components occupy space and must exchange signals. As the machine grows, communication becomes a primary physical limitation.

Suppose a device has characteristic diameter R. No signal can cross it faster than approximately

T_cross >= R/c,

where c is the speed of light.

A global operation requiring coordination across the entire device cannot be completed faster than this light-crossing time.

For small processors, communication delays may be negligible compared with gate times. For a machine distributed over planetary, stellar, or cosmic scales, signal propagation becomes dominant.

A Planck-scale spatial simulator with approximately 10^185 local registers would need an immense interconnection network. If the registers represented neighboring spatial cells, local communication might be arranged efficiently. Yet global constraints, long-range fields, error decoding, and synchronization could still require information to propagate across large portions of the machine.

The architecture would face two competing demands.

To fit many components into a small region, they must be packed densely. Dense packing increases unwanted interactions, thermal load, fabrication difficulty, and gravitational energy density.

To reduce cross-talk and energy density, components can be spread apart. This increases communication distance and latency.

The tradeoff can be summarized as

greater density -> stronger local interference and gravity,

greater separation -> slower coordination and larger apparatus.

Communication also consumes bandwidth. If each of N processors must exchange k bits or qubits per update, the total communication burden scales at least as

B_comm approximately kN

for local exchange, and may scale much faster for highly connected architectures.

Quantum communication adds further constraints. Unknown quantum states cannot be copied freely. Entanglement distribution requires channels, quantum memories, purification, swapping, or error correction. Loss and decoherence increase with distance and complexity.

A distributed quantum computer must preserve phase relationships and timing across its network. Clock synchronization is itself a relativistic and gravitational problem. Clocks at different gravitational potentials or velocities do not advance identically. For sufficiently large systems, spacetime geometry becomes part of the control architecture.

The finite speed of communication therefore prevents a universal instantaneous update. Any physically faithful computation must have a causal structure of its own.

This suggests that a realistic quantum spacetime compiler should be organized locally. It should route information only where dependencies require it and avoid global synchronization whenever possible.

Locality is not merely a property to be simulated. It is also a constraint on the simulator.

10.4 Decoherence and Fault Tolerance

Quantum computation depends on maintaining coherent relationships among states. Real systems continually interact with their environments, causing information to leak into uncontrolled degrees of freedom.

This process is called decoherence.

A coherent superposition such as

a|0> + b|1>

can become correlated with environmental states:

a|0>|E_0> + b|1>|E_1>.

If the environmental states become distinguishable and inaccessible, interference between the alternatives is suppressed from the perspective of the controlled system.

Large quantum systems are especially vulnerable because they present many channels for unwanted interaction. More qubits mean more surfaces, couplings, control lines, defects, thermal modes, and possible correlated errors.

The challenge is not simply to keep each qubit isolated. Computation also requires intentional interactions. The machine must be isolated from uncontrolled influences while remaining controllable through precisely engineered ones.

Quantum error correction addresses this difficulty by encoding logical information redundantly across larger physical systems. Errors are detected indirectly through syndrome measurements and corrected without directly measuring the protected logical state.

Fault tolerance requires the physical error rate to remain below a threshold determined by the code and noise model. Even below threshold, the overhead grows as the desired logical error rate decreases.

For a computation with G logically significant operations, a rough reliability target is

p_L G << 1,

where p_L is the logical error probability per operation.

If

G approximately 10^246,

the required aggregate reliability becomes extreme. It would be physically unreasonable to demand that one monolithic computation proceed across so many operations with no uncorrected failure.

A meaningful large-scale simulation would need a different concept of accuracy. It would have to tolerate local error, statistical uncertainty, and coarse-grained equivalence.

The distinction between exact microscopic fidelity and observable fidelity becomes important.

A simulation may contain many microscopic errors while still predicting a robust macroscopic observable accurately. Hydrodynamic quantities, phase classifications, conserved charges, or long-distance correlations may be insensitive to individual local mistakes.

A quantum spacetime compiler should therefore protect the information relevant to the query rather than every microscopic amplitude equally.

This may involve:

encoding only selected logical sectors,

using error mitigation for shallow components,

employing topological or bosonic protection,

preserving symmetries and conservation laws,

and reporting uncertainty in the final observable.

Decoherence also complicates claims about infinite-dimensional systems. High-energy oscillator states, closely spaced levels, and delicate nonclassical superpositions may decay or dephase rapidly. The number of levels that exist formally can be much larger than the number that remain coherent long enough to use.

Thus,

d_coherent <= d_accessible <= dim H.

The coherent operational dimension may be the quantity that actually limits quantum algorithms.

Fault tolerance can enlarge this dimension, but only by consuming additional physical states, energy, control bandwidth, and time.

10.5 Gravitational Back-Reaction

Any sufficiently powerful quantum computer becomes a gravitational system. Its energy, momentum, pressure, and information-bearing structures contribute to spacetime curvature.

For ordinary devices, this effect is negligible. For a machine approaching cosmic information densities, it becomes dominant.

A computer intended to simulate spacetime is not located outside spacetime. It is made of matter and fields within the geometry it seeks to model. As its mass-energy increases, the machine changes that geometry.

If energy E is concentrated within a radius R small enough that

R approximately less than or equal to 2GE/c^4,

the region approaches the condition associated with gravitational collapse.

The exact treatment depends on geometry and dynamics, but the lesson is clear: a finite region cannot contain unlimited energy or information without forming a horizon.

This creates a fundamental conflict for exhaustive simulation. To represent more independent states, the machine needs more physical degrees of freedom. Those degrees of freedom require

energy and structure. Concentrating them changes the spacetime and may eventually hide the machine behind a horizon.

A simulator dense enough to contain the supposed independent information of a Planck-scale cosmic raster would not remain a passive observer. It would become one of the most massive and geometrically significant objects in the Universe.

Its gravity could distort clocks, alter communication pathways, produce horizons, and change the very system being represented.

This feedback may be expressed as

computer state -> stress-energy -> spacetime curvature -> computer dynamics.

The simulator and simulated geometry become coupled.

A complete theory of computational limits in quantum gravity must therefore be self-consistent. It cannot specify qubit count and gate speed while assuming a fixed background geometry unaffected by the apparatus.

Gravitational back-reaction also restricts measurement. Attempting to resolve an extremely small region requires concentrating energy into that region. Beyond some scale, adding energy may create a black hole larger than the target resolution.

The attempt to inspect finer structure can therefore destroy access to that structure.

This suggests that the Planck scale is not merely a technological frontier. It may represent a domain in which the distinction between measuring apparatus, stored information, and spacetime geometry breaks down.

Hilbert harvesting at this scale cannot be discussed independently of gravity.

10.6 Self-Reference and the Simulator-Inside-the-Universe Problem

A final limitation arises from self-reference. A quantum computer capable of representing the entire observable Universe would itself be part of that Universe.

If the simulation is complete, it must include the computer, the computer's internal simulation, and the observer reading the output.

This produces a recursive structure:

Universe contains simulator,

simulator represents Universe,

represented Universe contains represented simulator,

represented simulator represents represented Universe,

and so on.

A finite physical machine cannot contain an independent, equally detailed copy of the entire system that includes the machine itself without some form of compression, abstraction, or self-reference.

The problem resembles a map required to contain a full-scale map of the territory, including the map itself. A perfectly detailed physical copy would require at least as many distinguishable states as the original system. If the copy is contained within the original, it cannot be physically separate and equally extensive without duplication of resources.

For a classical N-bit system, an arbitrary exact state requires N bits to represent. A simulator occupying some of those N bits cannot generally allocate N additional independent bits to a complete external copy of the whole system.

A quantum representation does not remove the difficulty. A faithful simulator must possess a Hilbert space capable of representing the relevant state space of the simulated system. If the simulated system includes the simulator, the required representation becomes self-referential.

One possible resolution is that the Universe simulates itself simply by evolving. Its physical state is not a stored model of reality; it is reality.

In this sense, the most exact simulation of the Universe is the Universe itself.

A constructed computer can represent only a projection, approximation, effective model, or selected observable. It may reproduce important dynamics without duplicating every degree of freedom.

Self-reference also limits prediction. A simulator may calculate a future outcome that is communicated to an agent whose behavior then changes because of the prediction. The prediction becomes part of the causal system it predicts.

In quantum mechanics, measurement choices and interventions further complicate the separation between model and modeled system. The simulator cannot always be treated as an external observer with unrestricted access to a fixed universal state.

The practical implication is that a quantum spacetime compiler must define its scope. It should specify:

which subsystem is modeled,

which degrees of freedom are treated as environment,

which observers and apparatus are included,

which observables are requested,

and what level of approximation is accepted.

The goal cannot be an independent, exhaustive, perfectly self-contained duplicate of all reality.

Finite energy restricts the number of accessible states. Finite time restricts how many transformations can occur. Finite size restricts communication. Decoherence restricts coherent control. Gravity restricts energy and information density. Self-reference restricts complete internal duplication.

Together, these limits define the true boundary of Hilbert-space harvesting.

The formal Hilbert space may remain infinite, but the operational harvest of any finite observer is bounded by the physics of the observer and apparatus.

This leads to a strengthened statement of the Hilbert Harvest Principle:

The useful state space of a physical quantum system is the subset that can be reached, controlled, preserved, connected, measured, and interpreted without violating finite energy, finite time, causality, fault tolerance, gravitational consistency, or self-reference.

The proper goal is therefore not to conquer the entire Hilbert space. It is to locate those structured regions of state space from which finite physical questions can be answered efficiently and reliably.

The limits described here do not make quantum spacetime computation meaningless. They define its realistic purpose.

A viable machine would not reproduce the Universe as an independent Planck-scale duplicate. It would construct bounded models, preserve relevant correlations, and extract observables within declared physical and epistemic limits.

Its success would be measured not by infinite capacity, but by disciplined selection of what finite observers can meaningfully know.

11. What Infinite Hilbert Spaces Can Productively Provide

11.1 Continuous-Variable Quantum Computation

Infinite-dimensional Hilbert spaces are most naturally realized in systems with continuous observables or unbounded excitation ladders. Examples include electromagnetic field modes, optical cavities, mechanical resonators, vibrational modes, and collective excitations in matter.

Instead of encoding information only in two discrete states, a continuous-variable system may use quantities analogous to position and momentum. For an electromagnetic mode, these are often represented by two field quadratures, x and p, satisfying a noncommutative relation of the form

[x, p] = i hbar.

A quantum state may then be described by a wavefunction over x, by a distribution in phase space, or by superpositions of oscillator number states.

The principal advantage is structural compatibility. Many physical systems are already described by continuous fields, oscillations, amplitudes, phases, and spectra. A continuous-variable quantum computer may therefore represent such systems more directly than a binary register.

A discretized qubit simulation must choose a finite grid, encode numerical amplitudes, manage truncation, and implement arithmetic through gates. An oscillator may embody the relevant variable natively.

This does not mean that an oscillator supplies infinite numerical precision. The usable range and resolution remain finite. Nevertheless, a finite-energy continuous-variable platform can provide an efficient representation of structured functions and bosonic dynamics.

Operations such as displacement, phase rotation, squeezing, interference, and mode coupling can transform these states directly. Some of these transformations correspond to simple physical interactions even when their classical simulation requires tracking large matrices or distributions.

Continuous-variable systems are therefore especially promising for:

bosonic field simulation,

wave propagation,

molecular vibration,

spectral estimation,

Gaussian transformations,

and hybrid analog-digital computation.

Their value lies in matching the architecture of the computer to the architecture of the problem.

11.2 Bosonic Quantum Error Correction

One of the most productive uses of an infinite-dimensional Hilbert space is not to store an unlimited number of raw symbols, but to protect a small number of logical quantum states.

A bosonic mode possesses many available number states:

|0>, |1>, |2>, ...

These states can be combined into structured codewords representing one logical qubit. The redundancy is distributed across the oscillator's larger state space.

Schematically, one may define

|0_L> = structured superposition of oscillator states,

|1_L> = another structured superposition.

The logical information is not identified with one fragile energy level. It is encoded across several levels or across a patterned region of phase space.

Different bosonic coding strategies protect against different errors. Some are designed to detect photon loss. Others protect against small displacements in position and momentum. Some exploit rotational symmetry, parity, or separated coherent-state components.

The important conceptual reversal is that the infinite ladder is used vertically rather than horizontally. Instead of treating each level as another independent memory symbol, the levels cooperate to stabilize one logical distinction.

This may reduce hardware overhead because one oscillator can provide internal redundancy that would otherwise require many separate two-level systems.

The operational Hilbert harvest is then not

many fragile bits per oscillator,

but

fewer, more reliable logical qubits per oscillator.

This is often the more valuable use of the available dimension.

The same idea may be relevant to spacetime information. If fundamental geometry or field modes possess large local state spaces, their physically useful role may be to encode robust relational or topological information rather than independent classical values at every location.

Redundancy within Hilbert space may be a mechanism of stability, not merely capacity.

11.3 Quantum Sensing and Metrology

Infinite-dimensional and continuous-variable systems are exceptionally valuable for measurement because their states respond sensitively to small physical perturbations.

A phase shift, displacement, frequency change, acceleration, force, or field can alter the state of an oscillator. Interference converts that alteration into a measurable signal.

The goal is not to extract every parameter of the quantum state. It is to estimate one selected physical quantity with high precision.

Suppose a parameter theta is encoded through a transformation

|psi(theta)> = U(theta)|psi_0>.

The experiment is designed so that small changes in theta produce distinguishable changes in measurement statistics.

Useful parameters may include:

time,

frequency,

magnetic field,

electric field,

acceleration,

rotation,

gravitational potential,

mechanical displacement,

temperature,

or molecular concentration.

Squeezed states are particularly important. In an ordinary minimum-uncertainty state, fluctuations are distributed between conjugate variables. Squeezing reduces uncertainty in one variable while increasing it in the other.

The uncertainty relation remains

Delta x Delta p >= hbar/2.

Squeezing does not remove quantum uncertainty. It reallocates it toward the variable that matters less for the chosen measurement.

If the experiment measures x, then reducing Delta x can improve sensitivity even though Delta p grows.

This principle is used in precision optics and interferometry. The infinite-dimensional state space supplies a broad family of nonclassical states whose fluctuations can be engineered for a particular task.

The harvest is therefore improved parameter sensitivity per unit energy, time, or probe particle.

A spacetime-oriented application could involve detecting tiny changes in distance, clock rate, gravitational field, curvature, or wave propagation. Such measurements would not image spacetime cell by cell. They would extract selected geometric observables with exceptional sensitivity.

This is a far more realistic use of Hilbert-space richness than exhaustive Planck-scale rendering.

11.4 Analog Quantum Simulation

A quantum system can sometimes simulate another quantum system by reproducing its relevant Hamiltonian, symmetry, mode structure, or interaction graph.

This is analog quantum simulation.

Instead of compiling every physical process into a long sequence of universal gates, the simulator is engineered so that its natural evolution parallels the target system.

Examples may involve:

photons simulating bosonic transport,

trapped particles simulating lattice models,

superconducting resonators simulating interacting modes,

or collective excitations simulating field dynamics.

The advantage comes from direct physical correspondence.

If the target system contains bosonic occupation numbers, an oscillator-based simulator may represent them without converting every occupation value into a binary string. If the target contains local couplings, the simulator may implement those couplings physically.

The simulator still requires truncation and calibration. Its Hilbert space may be infinite formally, but only a finite low-energy sector is controlled. The analog correspondence is also approximate. Imperfections in the simulator become errors in the modeled dynamics.

Nevertheless, analog simulation can make otherwise inaccessible quantum behavior observable.

The relevant measure is not whether the simulator reproduces every amplitude of the target system. It is whether it reproduces the observables, phases, correlations, and dynamical regimes of interest.

For quantum spacetime research, analog simulators might model:

horizon-like behavior,

field propagation in curved effective geometries,

particle production,

topological phases,

causal networks,

or entanglement patterns associated with emergent geometry.

Such systems would not be complete models of the Universe. They would be laboratory realizations of selected structural features.

The ability to isolate one mechanism and study it experimentally is a major scientific harvest from large Hilbert spaces.

11.5 Multimode Communication and Entanglement

A field contains many distinguishable modes. These may differ in frequency, polarization, time of arrival, spatial profile, propagation direction, or orbital angular momentum.

Each mode can carry information.

A communication channel need not therefore be restricted to one binary distinction per transmitted particle. A photon may occupy one of several modes, or a quantum state may be distributed coherently across many modes.

If d modes are reliably distinguishable, one ideal symbol can identify

log2 d

bits of classical labeling information.

For example, 256 modes could in principle label 8 bits per symbol. The practical capacity would be reduced by loss, detector errors, cross-talk, noise, and imperfect mode preparation.

Multiple modes can also be entangled. A state may take the form

|psi> = sum_n c_n |n>_A |n>_B,

where the systems A and B share correlated mode labels or occupation numbers.

High-dimensional entanglement can support:

larger alphabets,

more complex correlation structures,

improved resilience in selected noise environments,

distributed sensing,

and richer tests of quantum nonclassicality.

The useful resource is not infinity itself. It is the number of modes that can be reliably separated and controlled.

Mode multiplexing illustrates the Hilbert Harvest Principle particularly well. A single physical medium may possess a vast formal state space, but engineering determines how many channels remain distinguishable after propagation and detection.

The harvest is increased communication capacity within finite bandwidth, energy, and error constraints.

A future quantum network could combine discrete qubits, bosonic modes, and continuous variables. Different encodings would be used for memory, transmission, sensing, and error correction according to their strengths.

11.6 Structured Function Representation

A wavefunction is not merely a list of independent numbers. It is often a structured function generated by a small set of physical parameters and dynamical rules.

For example, a coherent state may be characterized by a complex amplitude alpha even though its expansion contains infinitely many number states:

|alpha> = exp(-|alpha|^2/2) sum_n [alpha^n/sqrt(n!)] |n>.

The state occupies an infinite-dimensional Hilbert space, but its description is compact because all amplitudes are generated from alpha.

Similarly, a Gaussian wavepacket can be specified through parameters such as center, width, phase, and momentum. It does not require an independent stored number at every point.

This reveals an important form of quantum compression:

large basis support does not necessarily imply large description complexity.

A physical system can embody a structured function directly. Its evolution may transform that function without explicitly calculating its value at every coordinate.

The system may then provide selected outputs such as:

samples from the distribution,

expectation values,

moments,

overlaps,

correlation functions,

or spectral components.

This is precisely the kind of capability required by the generative-spacetime framework.

A quantum state could implicitly represent a large family of possible geometries or histories. The observer would not reconstruct the entire function. The observer would ask for selected projections or observables.

The value of the representation depends on structure. If the encoded function is arbitrary, preparing and learning it may require resources proportional to its complexity. If it belongs to a compact family generated by lawful dynamics, a small number of parameters or operations may suffice.

The distinction parallels classical compression. A simple function can generate billions of samples, while a random table generally cannot be compressed.

Infinite-dimensional Hilbert spaces are therefore most useful when the relevant amplitudes are connected by symmetry, dynamics, analyticity, locality, or another generating rule.

They are least useful when treated as collections of unrelated numerical storage slots.

11.7 Query Advantage Rather Than Exhaustive Knowledge

The strongest practical conclusion is that quantum advantage should be understood as query advantage.

A quantum system may represent an enormous state implicitly. A useful algorithm arranges its dynamics so that one selected property becomes accessible through a limited measurement.

The advantage is not possession of exhaustive knowledge. It is efficient access to a particular answer.

A quantum device may estimate:

an energy,

a transition probability,

a phase,

a correlation,

a period,

a spectral gap,

or the expectation value of an observable.

It need not reveal every amplitude or every microscopic configuration.

This is the correct model for quantum spacetime computation.

Suppose a state implicitly represents many possible histories. A useful query might ask:

What is the probability that a horizon forms?

What is the expected curvature in a defined region?

Which boundary conditions are consistent with this observation?

What correlation links two detectors?

What effective geometry emerges at large scale?

The answer may require only a few classical numbers, even though the underlying state occupies a huge Hilbert space.

The gain comes from allowing the quantum system to perform interference and evolution internally while exporting only the relevant statistic.

This can be expressed schematically as

large implicit state + carefully chosen observable -> small useful output.

The strategy fails when the query demands exhaustive detail:

large implicit state + request for all components -> enormous output.

The No-Free-Readout Principle remains intact.

The most important scientific skill is therefore choosing the observable. A poorly framed question may require impossible reconstruction. A well-framed question may isolate a robust quantity that can be estimated without enumerating microscopic reality.

This suggests a hierarchy of increasingly demanding outputs:

one decision bit,

one estimated parameter,

one correlation function,

one sampled configuration,

one reduced state,

one local field map,

one global high-resolution state,

the full spacetime raster.

Quantum computation is most plausible near the top of this hierarchy only for very small systems. For cosmic-scale problems, useful work must remain near the lower-output end.

Infinite Hilbert spaces can productively provide flexible encodings, natural analog representations, protected logical states, high-dimensional correlations, and exceptional measurement sensitivity.

They cannot provide infinite usable information to a finite observer.

Their deepest value is generative. They allow a finite physical system to embody structured possibilities whose selected consequences can be harvested without exhaustive expansion.

The practical research program should therefore focus on three questions:

Which physical observables contain the most scientific value?

Which Hilbert-space structures encode those observables efficiently?

Which preparation and measurement protocols extract them with the least energy, time, and error?

These questions replace the impossible goal of occupying all formal states with the achievable goal of using the right states.

The future of quantum technology will not be determined by how close a machine comes to infinity. It will be determined by how intelligently it converts a finite accessible portion of Hilbert space into reliable answers.

12. Toward an Operational Theory of Computable Reality

12.1 The Universe Is Not a Movie Thesis

The central thesis of this paper is that the Universe should not be understood computationally as a prewritten movie composed of independently stored Planck-scale frames. A four-dimensional raster with approximately 10^246 cells is a possible representation of cosmic history, but it is not necessarily the physical architecture of reality.

The raster count measures how many entries would be required if every spatial location and every temporal interval were explicitly sampled. It does not establish that each entry corresponds to an independent degree of freedom, a distinct memory location, or an irreducible physical fact.

Physical states are constrained by laws, symmetries, conservation principles, gauge relationships, causal structure, and quantum correlations. Successive moments are generated from preceding conditions rather than selected arbitrarily. Nearby regions are correlated. Coordinate descriptions contain redundancy. Gravitational entropy bounds suggest that the independent information associated with a bounded region may scale more nearly with boundary area than with interior volume.

The complete Planck-scale movie is therefore better interpreted as a maximally expanded classical output than as the minimum description of the Universe.

This distinction can be summarized as

expanded spacetime raster != independent physical information.

A physical history may instead be represented by

state + dynamics + constraints + boundary conditions.

The history generated from this description can be enormous even when the description itself is comparatively compact.

The Universe Is Not a Movie thesis does not require denial of four-dimensional spacetime, relativistic block descriptions, or the physical reality of past events. It claims only that a four-dimensional geometry need not function as a literal storage array in which every spacetime point holds an independent record.

A film is constructed from separate frames. Physical history may be generated through lawful relationships among states.

The more appropriate computational metaphor is therefore not a movie archive but a quantum process capable of producing finite observable consequences.

12.2 Formal Infinity and Finite Observation

Infinite-dimensional Hilbert spaces are not inconsistent with finite physical information. The apparent contradiction arises only when formal state-space dimension is confused with operational accessibility.

A harmonic oscillator possesses the basis

|0>, |1>, |2>, ...

without a highest formal occupation number. A quantum field contains infinitely many idealized modes in a continuum description. A wavefunction can involve continuous variables and complex amplitudes specified mathematically with arbitrary precision.

Yet a finite observer operates with finite energy, finite time, finite apparatus, finite resolution, and nonzero error.

The operationally usable state space is therefore restricted.

The relevant quantity is not simply

dim H,

but

d_op(E, R, T, epsilon),

the number of states that can be prepared, controlled, preserved, distinguished, and measured under specified physical limits.

The corresponding operational information capacity is

H_op = log2 d_op.

For any finite experimental implementation, H_op remains finite even when dim H is infinite.

This resolves the original question of whether Hilbert-space infinity can provide unlimited memory. The infinity guarantees that the mathematical theory does not impose a fixed finite ceiling on its basis. It does not guarantee that a finite device can reach every basis state or distinguish arbitrarily close amplitudes.

Formal infinity provides possibility without universal access.

Finite observation also means that reality is encountered through projections. An observer does not receive the full quantum state of a system. The observer receives measurement outcomes associated with selected observables.

A quantum state may support an enormous range of possible questions. Each experiment extracts only a limited answer.

This creates a hierarchy:

formal state space,

physically accessible subspace,

controllable code space,

measurable observables,

recorded classical outcomes.

Each level is generally smaller than the one above it.

The infinity of Hilbert space is therefore physically useful not because it gives finite observers infinite knowledge, but because it supplies a rich substrate from which useful finite structures can be selected.

12.3 The Generative Spacetime Principle

The Generative Spacetime Principle may be stated as follows:

A physical spacetime history should be represented by the smallest set of states, laws, constraints, causal relationships, and reconstruction procedures sufficient to generate its observable consequences, rather than by an independently stored value at every spacetime point.

This principle separates ontology from representation. It does not insist that the Universe is literally a computer or that spacetime is produced by software. It states that an adequate physical description should exploit generative structure wherever that structure exists.

A generative formulation may be written schematically as

$H = F(\psi_0, L, C, B)$,

where

H = physical history,

ψ_0 = initial or boundary state,

L = dynamical laws,

C = constraints,

B = boundary conditions.

A finite observer does not ordinarily require the complete H. The observer requests an answer associated with an observable O:

$A_O = Q(H, O)$.

A quantum spacetime compiler would attempt to calculate A_O without constructing every component of H explicitly.

The principle has several implications.

First, time is primarily transformation rather than storage. Later states arise from earlier states through dynamics.

Second, coordinates are not automatically physical memory addresses. Coordinate systems may contain gauge and representational redundancy.

Third, local microscopic detail should be generated only when required by a physically meaningful observable.

Fourth, the computational cost of a question should depend on its causal support, resolution, precision, and output size.

Fifth, a complete description of physically accessible reality may be far smaller than a complete coordinate raster, while still being difficult to decode.

Generative spacetime therefore unites the Hilbert Harvest Principle and the No-Free-Readout Principle.

Hilbert-space structure provides an implicit representation.

Generative dynamics transform that representation.

Measurement extracts a bounded answer.

The output remains finite because the observer and recording apparatus are finite.

12.4 Research Questions

The framework developed here raises a set of research questions that cross quantum information, quantum gravity, computation, cosmology, and measurement theory.

One major question concerns the operational Hilbert dimension. How should d_op(E, R, T, epsilon) be defined for realistic quantum systems with noise, gravity, control restrictions, and finite measurement resources?

A second concerns gravitational information. Under what conditions can horizon-area entropy estimates be interpreted as bounds on operationally distinguishable quantum states? How do these bounds apply to cosmological regions that are not ordinary isolated black holes?

A third concerns spacetime emergence. If geometry arises from entanglement, algebraic relationships, causal order, tensor-network structure, or boundary information, which quantities replace the naive Planck voxel as the elementary computational object?

A fourth concerns reconstruction complexity. Even if an interior spacetime is compactly encoded, how difficult is it to reconstruct a selected local observable? Which interior questions are efficiently decodable, and which require computational effort comparable to exhaustive evolution?

A fifth concerns adaptive resolution. Can a quantum simulation refine spacetime only where an observable is sensitive to microscopic structure while maintaining a provable error bound elsewhere?

A sixth concerns quantum error correction. Are gravitational or holographic encodings naturally analogous to error-correcting codes? Can bulk information remain robust because it is redundantly represented in boundary or relational degrees of freedom?

A seventh concerns causal computation. How should quantum algorithms be organized when the computational dependency graph must itself respect relativistic causality and dynamical geometry?

An eighth concerns output-sensitive complexity. Can a general theory classify physical questions according to the size of their answers, causal support, required precision, and reconstruction cost?

A ninth concerns self-reference. What is the strongest form of universal prediction available to a computer that is itself embedded within the system being predicted?

A tenth concerns empirical testing. Which observable signatures could distinguish a fundamentally local Planck-cell ontology from a holographic, relational, or emergent description of spacetime?

These questions suggest that the computational study of reality should focus less on raw state-space size and more on accessibility, encoding, decoding, and observation.

12.5 Criteria for a Quantum Spacetime Compiler

A credible quantum spacetime compiler would require more than a large quantum processor. It would need a disciplined theoretical and computational framework.

The first criterion is physical validity. The encoded state and evolution rules must correspond to a defined physical theory or effective model. The compiler cannot hide unknown physics behind computational terminology.

The second criterion is causal consistency. The calculation must respect the causal structure of the modeled system, including finite signal propagation and the distinction between correlations and controllable influences.

The third criterion is gauge consistency. Outputs should be expressed through gauge-invariant or relational quantities whenever possible. Coordinate-dependent artifacts should not be mistaken for physical predictions.

The fourth criterion is resource accounting. Every calculation should report logical qubits, physical overhead, circuit depth, communication demands, energy assumptions, and readout cost.

The fifth criterion is resolution transparency. The compiler should state where fine-grained representations were used, where coarse-graining was introduced, and how unresolved structure could affect the answer.

The sixth criterion is uncertainty reporting. Numerical, statistical, model, truncation, and state-preparation uncertainties should be separated rather than combined into an unexplained confidence value.

The seventh criterion is gravitational self-consistency. At sufficiently large scales or energy densities, the physical computer and its control apparatus cannot be assumed to have negligible gravitational effects.

The eighth criterion is observable orientation. The computation should be designed around a defined physical question rather than around universal state reconstruction.

The ninth criterion is reproducibility. Another investigator using the same state specification, laws, approximations, and query should be able to reproduce the reported result within the declared uncertainty.

The tenth criterion is falsifiability. The compiler should produce predictions that can be compared with observation, experiment, or independent theoretical calculation.

A machine satisfying these criteria would not be an omniscient simulator. It would be a controlled scientific instrument for translating quantum models into bounded predictions.

12.6 Philosophical Implications

The framework has implications beyond engineering.

One implication concerns the meaning of physical information. Information is not merely a number attached to an abstract mathematical space. It requires distinguishable alternatives, physical encoding, possible transformations, and some form of observation or record.

Another concerns the status of the wavefunction. A quantum state may contain the structure needed to predict many possible measurements, yet it is not equivalent to a classical catalog containing the outcomes of all those measurements simultaneously.

A third concerns spacetime. If physically independent information scales with boundaries, causal relationships, or entanglement structure, then spacetime geometry may be a macroscopic expression of deeper informational organization.

A fourth concerns observation. Finite observers never encounter the complete Universe as a total object. They encounter causal, coarse-grained, measurement-dependent projections.

This does not imply that reality is subjective. Different observers can compare records, test predictions, and identify invariant relationships. It means that physical knowledge is operationally mediated.

A fifth implication concerns determinism. Even if a compact law and initial state generate a history, prediction may remain limited by chaos, quantum uncertainty, computational irreducibility, unavailable boundary data, or self-reference.

A sixth concerns simulation. To say that the Universe can be described generatively is not to say that it is artificial or that it runs on an external computer. Physical law itself may be the generator.

A seventh concerns completeness. The most complete scientifically meaningful account may not be the longest possible list of microscopic values. It may be the smallest framework capable

of producing all observable regularities and correctly answering all physically available questions.

This reverses a common assumption. More data do not always mean more understanding. An exhaustive raster may contain vastly more entries than a generative theory while explaining less.

12.7 Final Synthesis

The original estimate of approximately 10^246 Planck-scale spacetime cells revealed the impossibility of treating the complete observable cosmic history as an ordinary uncompressed database.

A full four-dimensional memory would require approximately 10^246 logical binary degrees of freedom for only one bit per cell. A sequential simulation could reduce active memory to approximately 10^185 spatial registers but would retain an enormous history of local updates. Massive parallelism would exchange runtime for hardware, communication, energy, and gravitational cost.

Infinite-dimensional Hilbert space does not remove these burdens. It contains unlimited formal basis states, but finite energy, time, precision, coherence, and measurement restrict the accessible portion.

The operational quantity is

$H_{op} = \log_2 d_{op}$,

not the formal dimension alone.

An 818-qubit register can address approximately 10^246 alternatives, but it cannot store one independent bit at every address. A superposition can transform many amplitudes coherently, but one measurement does not reveal them all. A compact quantum representation can answer selected questions, but it cannot print an exhaustive classical movie without paying the output cost.

These conclusions yield three principles.

The Hilbert Harvest Principle states that usable quantum capacity is determined by controllable and distinguishable states under finite constraints.

The No-Free-Readout Principle states that N independently requested classical bits require an output process and record capable of supporting N classical bits.

The Generative Spacetime Principle states that physical history should be represented through states, laws, constraints, boundaries, and queries rather than through an independent stored value at every spacetime point.

Together, these principles replace the goal of exhaustive reality capture with a more meaningful objective:

identify the smallest physically valid quantum representation from which finite observers can obtain every operationally accessible answer.

The Universe may contain an immense or formally infinite space of possible states. A finite observer cannot occupy, inspect, or export that infinity. The observer can prepare a bounded state, apply a bounded transformation, select an observable, and record a bounded result.

The deepest resource is therefore not infinite storage.

It is the ability to ask the right finite question of an immense generative reality.

13. References

1. National Institute of Standards and Technology. 2022 CODATA recommended values of the fundamental physical constants: Planck length and Planck time. NIST Reference on Constants, Units, and Uncertainty, Gaithersburg, Maryland, 2024.
2. Planck Collaboration. Planck 2018 results. VI. Cosmological parameters. Astronomy and Astrophysics 641, A6, 2020. DOI: 10.1051/0004-6361/201833910.
3. Davis, T. M., and Lineweaver, C. H. Expanding confusion: Common misconceptions of cosmological horizons and the superluminal expansion of the Universe. Publications of the Astronomical Society of Australia 21, 97–109, 2004. DOI: 10.1071/AS03040.
4. Dirac, P. A. M. The Principles of Quantum Mechanics. Oxford University Press, Oxford, 1930.
5. von Neumann, J. Mathematical Foundations of Quantum Mechanics. Princeton University Press, Princeton, 1955. Original German edition published in 1932.
6. Feynman, R. P. Simulating physics with computers. International Journal of Theoretical Physics 21, 467–488, 1982. DOI: 10.1007/BF02650179.
7. Deutsch, D. Quantum theory, the Church-Turing principle and the universal quantum computer. Proceedings of the Royal Society of London A 400, 97–117, 1985. DOI: 10.1098/rspa.1985.0070.
8. Lloyd, S. Universal quantum simulators. Science 273, 1073–1078, 1996. DOI: 10.1126/science.273.5278.1073.
9. Lloyd, S. Ultimate physical limits to computation. Nature 406, 1047–1054, 2000. DOI: 10.1038/35023282.
10. Lloyd, S. Computational capacity of the Universe. Physical Review Letters 88, 237901, 2002. DOI: 10.1103/PhysRevLett.88.237901.
11. Nielsen, M. A., and Chuang, I. L. Quantum Computation and Quantum Information. Tenth Anniversary Edition. Cambridge University Press, Cambridge, 2010.
12. Holevo, A. S. Bounds for the quantity of information transmitted by a quantum communication channel. Problems of Information Transmission 9, 177–183, 1973.
13. Schumacher, B. Quantum coding. Physical Review A 51, 2738–2747, 1995. DOI: 10.1103/PhysRevA.51.2738.

14. Helstrom, C. W. Quantum Detection and Estimation Theory. Academic Press, New York, 1976.

15. Wootters, W. K., and Zurek, W. H. A single quantum cannot be cloned. Nature 299, 802–803, 1982. DOI: 10.1038/299802a0.

16. Dieks, D. Communication by EPR devices. Physics Letters A 92, 271–272, 1982. DOI: 10.1016/0375-9601(82)90084-6.

17. Landauer, R. Irreversibility and heat generation in the computing process. IBM Journal of Research and Development 5, 183–191, 1961. DOI: 10.1147/rd.53.0183.

18. Bennett, C. H. Logical reversibility of computation. IBM Journal of Research and Development 17, 525–532, 1973. DOI: 10.1147/rd.176.0525.

19. Bennett, C. H. The thermodynamics of computation: A review. International Journal of Theoretical Physics 21, 905–940, 1982. DOI: 10.1007/BF02084158.

20. Shor, P. W. Scheme for reducing decoherence in quantum computer memory. Physical Review A 52, R2493–R2496, 1995. DOI: 10.1103/PhysRevA.52.R2493.

21. Steane, A. M. Error correcting codes in quantum theory. Physical Review Letters 77, 793–797, 1996. DOI: 10.1103/PhysRevLett.77.793.

22. Knill, E., Laflamme, R., and Zurek, W. H. Resilient quantum computation: Error models and thresholds. Proceedings of the Royal Society of London A 454, 365–384, 1998. DOI: 10.1098/rspa.1998.0166.

23. Gottesman, D. Stabilizer Codes and Quantum Error Correction. Ph.D. dissertation, California Institute of Technology, Pasadena, 1997.

24. Terhal, B. M. Quantum error correction for quantum memories. Reviews of Modern Physics 87, 307–346, 2015. DOI: 10.1103/RevModPhys.87.307.

25. Braunstein, S. L., and van Loock, P. Quantum information with continuous variables. Reviews of Modern Physics 77, 513–577, 2005. DOI: 10.1103/RevModPhys.77.513.

26. Lloyd, S., and Braunstein, S. L. Quantum computation over continuous variables. Physical Review Letters 82, 1784–1787, 1999. DOI: 10.1103/PhysRevLett.82.1784.

27. Gottesman, D., Kitaev, A., and Preskill, J. Encoding a qubit in an oscillator. Physical Review A 64, 012310, 2001. DOI: 10.1103/PhysRevA.64.012310.

28. Mirrahimi, M., Leghtas, Z., Albert, V. V., Touzard, S., Schoelkopf, R. J., Jiang, L., and Devoret, M. H. Dynamically protected cat-qubits: A new paradigm for universal quantum

computation. New Journal of Physics 16, 045014, 2014. DOI: 10.1088/1367-2630/16/4/045014.

29. Michael, M. H., Silveri, M., Brierley, R. T., Albert, V. V., Salmilehto, J., Jiang, L., and Girvin, S. M. New class of quantum error-correcting codes for a bosonic mode. Physical Review X 6, 031006, 2016. DOI: 10.1103/PhysRevX.6.031006.

30. Lvovsky, A. I., and Raymer, M. G. Continuous-variable optical quantum-state tomography. Reviews of Modern Physics 81, 299–332, 2009. DOI: 10.1103/RevModPhys.81.299.

31. Paris, M. G. A., and Rehacek, J., editors. Quantum State Estimation. Lecture Notes in Physics, Volume 649. Springer, Berlin, 2004. DOI: 10.1007/b98673.

32. Zurek, W. H. Decoherence, einselection, and the quantum origins of the classical. Reviews of Modern Physics 75, 715–775, 2003. DOI: 10.1103/RevModPhys.75.715.

33. Schlosshauer, M. Decoherence, the measurement problem, and interpretations of quantum mechanics. Reviews of Modern Physics 76, 1267–1305, 2005. DOI: 10.1103/RevModPhys.76.1267.

34. Mandelstam, L., and Tamm, I. The uncertainty relation between energy and time in nonrelativistic quantum mechanics. Journal of Physics of the USSR 9, 249–254, 1945.

35. Margolus, N., and Levitin, L. B. The maximum speed of dynamical evolution. Physica D 120, 188–195, 1998. DOI: 10.1016/S0167-2789(98)00054-2.

36. Bremermann, H. J. Quantum noise and information. Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 4, 15–20, University of California Press, Berkeley, 1967.

37. Kolmogorov, A. N. Three approaches to the quantitative definition of information. Problems of Information Transmission 1, 1–7, 1965.

38. Chaitin, G. J. On the length of programs for computing finite binary sequences. Journal of the ACM 13, 547–569, 1966. DOI: 10.1145/321356.321363.

39. Zurek, W. H. Algorithmic randomness and physical entropy. Physical Review A 40, 4731–4751, 1989. DOI: 10.1103/PhysRevA.40.4731.

40. Bekenstein, J. D. Black holes and entropy. Physical Review D 7, 2333–2346, 1973. DOI: 10.1103/PhysRevD.7.2333.

41. Hawking, S. W. Particle creation by black holes. Communications in Mathematical Physics 43, 199–220, 1975. DOI: 10.1007/BF02345020.

42. Bekenstein, J. D. Universal upper bound on the entropy-to-energy ratio for bounded systems. Physical Review D 23, 287–298, 1981. DOI: 10.1103/PhysRevD.23.287.

43. Casini, H. Relative entropy and the Bekenstein bound. Classical and Quantum Gravity 25, 205021, 2008. DOI: 10.1088/0264-9381/25/20/205021.

44. Gibbons, G. W., and Hawking, S. W. Cosmological event horizons, thermodynamics, and particle creation. Physical Review D 15, 2738–2751, 1977. DOI: 10.1103/PhysRevD.15.2738.

45. Egan, C. A., and Lineweaver, C. H. A larger estimate of the entropy of the Universe. Astrophysical Journal 710, 1825–1834, 2010. DOI: 10.1088/0004-637X/710/2/1825.

46. ’t Hooft, G. Dimensional reduction in quantum gravity. In Salamfestschrift: A Collection of Talks, edited by A. Ali, J. Ellis, and S. Randjbar-Daemi, World Scientific, Singapore, 1993. arXiv: gr-qc/9310026.

47. Susskind, L. The world as a hologram. Journal of Mathematical Physics 36, 6377–6396, 1995. DOI: 10.1063/1.531249.

48. Bousso, R. A covariant entropy conjecture. Journal of High Energy Physics 1999, 004, 1999. DOI: 10.1088/1126-6708/1999/07/004.

49. Bousso, R. The holographic principle. Reviews of Modern Physics 74, 825–874, 2002. DOI: 10.1103/RevModPhys.74.825.

50. Bousso, R., Casini, H., Fisher, Z., and Maldacena, J. Proof of a quantum Bousso bound. Physical Review D 90, 044002, 2014. DOI: 10.1103/PhysRevD.90.044002.

51. Maldacena, J. M. The large N limit of superconformal field theories and supergravity. Advances in Theoretical and Mathematical Physics 2, 231–252, 1998. DOI: 10.4310/ATMP.1998.v2.n2.a1.

52. Gubser, S. S., Klebanov, I. R., and Polyakov, A. M. Gauge theory correlators from non-critical string theory. Physics Letters B 428, 105–114, 1998. DOI: 10.1016/S0370-2693(98)00377-3.

53. Witten, E. Anti-de Sitter space and holography. Advances in Theoretical and Mathematical Physics 2, 253–291, 1998. DOI: 10.4310/ATMP.1998.v2.n2.a2.

54. Ryu, S., and Takayanagi, T. Holographic derivation of entanglement entropy from the anti-de Sitter space/conformal field theory correspondence. Physical Review Letters 96, 181602, 2006. DOI: 10.1103/PhysRevLett.96.181602.

55. Jacobson, T. Thermodynamics of spacetime: The Einstein equation of state. Physical Review Letters 75, 1260–1263, 1995. DOI: 10.1103/PhysRevLett.75.1260.

56. Van Raamsdonk, M. Building up spacetime with quantum entanglement. General Relativity and Gravitation 42, 2323–2329, 2010. DOI: 10.1007/s10714-010-1034-0.

57. Vidal, G. Entanglement renormalization. Physical Review Letters 99, 220405, 2007. DOI: 10.1103/PhysRevLett.99.220405.

58. Swingle, B. Entanglement renormalization and holography. Physical Review D 86, 065007, 2012. DOI: 10.1103/PhysRevD.86.065007.

59. Almheiri, A., Dong, X., and Harlow, D. Bulk locality and quantum error correction in AdS/CFT. Journal of High Energy Physics 2015, 163, 2015. DOI: 10.1007/JHEP04(2015)163.

60. Pastawski, F., Yoshida, B., Harlow, D., and Preskill, J. Holographic quantum error-correcting codes: Toy models for the bulk/boundary correspondence. Journal of High Energy Physics 2015, 149, 2015. DOI: 10.1007/JHEP06(2015)149.

61. Harlow, D. The Ryu-Takayanagi formula from quantum error correction. Communications in Mathematical Physics 354, 865–912, 2017. DOI: 10.1007/s00220-017-2904-z.

62. Harlow, D. TASI lectures on the emergence of bulk physics in AdS/CFT. Proceedings of Science, TASI2017, 002, 2018. DOI: 10.22323/1.305.0002.

63. Bombelli, L., Lee, J., Meyer, D., and Sorkin, R. D. Space-time as a causal set. Physical Review Letters 59, 521–524, 1987. DOI: 10.1103/PhysRevLett.59.521.

64. Sorkin, R. D. Causal sets: Discrete gravity. In Lectures on Quantum Gravity, edited by A. Gomberoff and D. Marolf, 305–327, Springer, New York, 2005. DOI: 10.1007/0-387-24992-3_10.

65. Rovelli, C., and Smolin, L. Discreteness of area and volume in quantum gravity. Nuclear Physics B 442, 593–622, 1995. DOI: 10.1016/0550-3213(95)00150-Q.

66. Ashtekar, A., and Lewandowski, J. Quantum theory of geometry. I. Area operators. Classical and Quantum Gravity 14, A55–A82, 1997. DOI: 10.1088/0264-9381/14/1A/006.

67. Wheeler, J. A. Information, physics, quantum: The search for links. In Complexity, Entropy, and the Physics of Information, edited by W. H. Zurek, 3–28, Addison-Wesley, Redwood City, California, 1990.

68. Wheeler, J. A., and Ford, K. Geons, Black Holes, and Quantum Foam: A Life in Physics. W. W. Norton, New York, 1998.

69. Aaronson, S. The limits of quantum computers. Scientific American 298, 62–69, 2008. DOI: 10.1038/scientificamerican0308-62.

70. Preskill, J. Quantum computing in the NISQ era and beyond. Quantum 2, 79, 2018. DOI: 10.22331/q-2018-08-06-79.

71. Georgescu, I. M., Ashhab, S., and Nori, F. Quantum simulation. Reviews of Modern Physics 86, 153–185, 2014. DOI: 10.1103/RevModPhys.86.153.

72. Cirac, J. I., and Zoller, P. Goals and opportunities in quantum simulation. Nature Physics 8, 264–266, 2012. DOI: 10.1038/nphys2275.

73. Berry, D. W., Ahokas, G., Cleve, R., and Sanders, B. C. Efficient quantum algorithms for simulating sparse Hamiltonians. Communications in Mathematical Physics 270, 359–371, 2007. DOI: 10.1007/s00220-006-0150-x.

74. Low, G. H., and Chuang, I. L. Hamiltonian simulation by qubitization. Quantum 3, 163, 2019. DOI: 10.22331/q-2019-07-12-163.

75. Nielsen, M. A. A geometric approach to quantum circuit lower bounds. Quantum Information and Computation 6, 213–262, 2006.

76. Eisert, J., Cramer, M., and Plenio, M. B. Area laws for the entanglement entropy. Reviews of Modern Physics 82, 277–306, 2010. DOI: 10.1103/RevModPhys.82.277.

77. Orus, R. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. Annals of Physics 349, 117–158, 2014. DOI: 10.1016/j.aop.2014.06.013.

78. Biamonte, J., and Bergholm, V. Tensor networks in a nutshell. arXiv: 1708.00006, 2017.

79. Hayden, P., Nezami, S., Qi, X. L., Thomas, N., Walter, M., and Yang, Z. Holographic duality from random tensor networks. Journal of High Energy Physics 2016, 9, 2016. DOI: 10.1007/JHEP11(2016)009.

80. Cao, C., Carroll, S. M., and Michalakis, S. Space from Hilbert space: Recovering geometry from bulk entanglement. Physical Review D 95, 024031, 2017. DOI: 10.1103/PhysRevD.95.024031.

The author has published over 5,600 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.